

Canonicity results for mu-calculi: an algorithmic approach

Willem Conradie

Andrew Craig

wconradie@uj.ac.za
University of Johannesburg
South Africa

acraig@uj.ac.za
University of Johannesburg
South Africa

August 28, 2014

Abstract

We investigate the canonicity of inequalities of the intuitionistic mu-calculus. The notion of canonicity in the presence of fixed point operators is not entirely straightforward. In the algebraic setting of canonical extensions we examine both the usual notion of canonicity and what we will call tame canonicity. This latter concept has previously been investigated for the classical mu-calculus by Bezhanishvili and Hodkinson. Our approach is in the spirit of Sahlqvist theory. That is, we identify syntactically-defined classes of inequalities, namely the restricted inductive and tame inductive inequalities, which are, respectively, canonical or tame canonical. Our approach is to use an algorithm which processes inequalities with the aim of eliminating propositional variables. The algorithm we introduce is closely related to the algorithms ALBA and mu-ALBA studied by Conradie, Palmigiano, et al. It is based on a calculus of rewrite rules, the soundness of which rests upon the way in which algebras embed into their canonical extensions and the order-theoretic properties of the latter. We show that the algorithm succeeds on every restricted inductive inequality by means of a so-called proper run, and that this is sufficient to guarantee their canonicity. Likewise, we are able to show that the algorithm succeeds on every tame inductive inequality by means of a so-called tame run. In turn, this guarantees their tame canonicity.

Keywords: modal mu-calculus, Sahlqvist theory, canonical extension, ALBA algorithm, canonicity.

1 Introduction

The modal mu-calculus was defined in 1983 by Kozen [11] and is obtained by adding the least and greatest fixed point operators to the basic modal logic. An overview of the modal mu-calculus can be found in the chapter by Bradfield and Stirling [4]. Canonical models and completeness results for the finitary and infinitary logics defined by Kozen were obtained by Ambler et al. [1]. In our algebraic canonicity proofs we will build on the definitions of modal mu-algebras from [1]. The correspondence and completeness of logics with fixed point operators has been the subject of recent studies by Bezhanishvili and Hodkinson [3] and Conradie et al. [5]. Both of these works aim to develop a Sahlqvist-like theory for their respective fixed point settings.

Sahlqvist theory, first developed in 1975 [12], is one of the most important and powerful ideas in the study of modal and related logics. The theory consists of two parts: *canonicity* and *correspondence*. The Sahlqvist formulas are a recursively defined class of modal formulas with a particular syntactic shape. Any modal logic axiomatized by Sahlqvist formulas is strongly complete (via canonicity) with respect to its class of Kripke frames, and the latter is moreover guaranteed to be an elementary class. This last fact, that the class of frames can be characterized by first-order conditions, is the correspondence aspect of Sahlqvist theory. This Sahlqvist-style approach of describing a class of formulas of a certain syntactic shape for which correspondence and completeness results can be proved, has been extensively developed by van Benthem [13, 14] and others. The

so-called Sahlqvist–van Benthem algorithm is used to find the first-order condition that corresponds to a given Sahlqvist formula. Various generalizations of the Sahlqvist class exist, including the inductive formulas, introduced by Goranko and Vakarelov [10].

The work in [3] looks at both correspondence and a certain type of canonicity for the classical mu-calculus. They define *Sahlqvist fixed point formulas*: a syntactic class which allows for limited use of fixed point operators. A modified version of canonicity is proved for these Sahlqvist fixed point formulas. (We also note the related algebraic work on preservation of Sahlqvist fixed point equations under MacNeille completions by Bezhanishvili and Hodkinson [2].)

In contrast to the work on Sahlqvist fixed point formulas in [3], [5] examines correspondence only, and, using an algorithmic approach, obtains results for a broader class of formulas in the setting of bi-intuitionistic mu-calculus. The correspondence results are achieved by examining validity on the complex algebras dual to Kripke frames. The algorithmic approach of [5] builds on work by Conradie and Palmigiano [6] on canonicity and correspondence for distributive modal logic.

In this paper we will make use of purely algebraic and order-theoretic techniques. This approach has proved fruitful in obtaining completeness results for many non-classical logics.

We prove two different canonicity results for two classes of intuitionistic mu-formulas. We show that the members of a certain class of intuitionistic mu-formulas are canonical, in the sense of [3]; that is, they are preserved under certain modified canonical extensions. We refer to this modified form of canonicity (described in detail in Section 3) as *tame canonicity*. We also define a second class of formulas for which the usual notion of canonicity holds. However, in this second case we do not get completeness of the logics defined by the canonical mu-formulas. This lack of completeness is explained in more detail in Section 3.

Our methods use a variation of the algorithm ALBA (Ackermann Lemma Based Algorithm) developed in [6]. The key step in our algorithm μ^* -ALBA is the elimination of propositional variables via an Ackermann-style rule. We define *tame* and *proper* runs of our algorithm and show that all mu-inequalities that can be successfully processed by these runs are, respectively, tame canonical or canonical.

In Section 2 we define the languages in which we operate and establish the algebraic setting for the interpretation of these languages. Section 3 summarizes our canonicity results and outlines the method for achieving these results. The syntactic classes of inequalities for which we obtain our canonicity results are defined in Section 4; the examples in this section will assist the reader in getting to grips with the rather technical syntactic definitions. The algorithm μ^* -ALBA is presented in Section 5. The soundness of most of the rules of μ^* -ALBA follows easily from properties of the interpreting algebras. However, the proofs of the fixed point approximation rules and the Ackermann rules require more intricate algebraic manipulations. These proofs are given in Section 7 and Section 8 respectively. Some technical lemmas required for the proofs in Section 8 are given in Appendix 11. Section 6 describes a syntactically defined class of formulas, the term functions of which satisfy the order-theoretic conditions required by the fixed point approximation rules.

Section 9 demonstrates the tame canonicity of mu-inequalities on which a tame run of our algorithm succeeds and also shows the canonicity of mu-inequalities on which a proper run of μ^* -ALBA succeeds. In Section 10 we prove that the members of the two different syntactic classes defined in Section 4 are, respectively, tame canonical and canonical. To end the paper, in Section 11 we present two examples of the algorithm at work.

2 Language and interpretation

In this section we collect the essential details of the syntax and semantics we will be using. We have opted to work in an intuitionistic rather than classical setting for two reasons: firstly, it allows us to carefully disentangle the order theoretic properties of connectives which make our approach tick, in a way that would seem unnecessary and pedantic if, e.g., classical negation was available; secondly, this added generality comes at very little extra cost. Since the focus of this paper is canonicity, we will work almost exclusively with algebraic semantics, which in this case takes the form of bi-Heyting algebras with additional modal operators. The relational

semantics can be given, as usual, by intuitionistic Kripke frames with additional relations for interpreting the modalities, see e.g., [15] and [7].

Modal bi-Heyting algebras. A *bi-Heyting algebra* is an algebra $\mathbf{A} = (A, \wedge, \vee, \rightarrow, -, \top, \perp)$ such that both the reducts $(A, \wedge, \vee, \rightarrow, \top, \perp)$ and $(A, \wedge, \vee, -, \top, \perp)^\theta$ are Heyting algebras. A *modal bi-Heyting algebra* is an algebra

$$\mathbf{A} = (A, \wedge, \vee, \rightarrow, -, \top, \perp, \Box, \Diamond)$$

such that $(A, \wedge, \vee, \rightarrow, -, \top, \perp)$ is a bi-Heyting algebra and \Box and \Diamond preserve finite meets and joins, respectively. We observe that \rightarrow and $-$ satisfy the inequalities

$$u \wedge v \leq w \quad \text{iff} \quad u \leq v \rightarrow w \quad \text{and} \quad u - v \leq w \quad \text{iff} \quad u \leq v \vee w.$$

The completely join-irreducible elements and completely meet-irreducible elements of a complete lattice will play a very important role in our algorithmic approach.

Definition 2.1. Let C be a complete lattice. Then

- (i) $j \in C$ is *completely join-irreducible* if for any $X \subseteq C$, if $j = \bigvee X$ then $j = x$ for some $x \in X$;
- (ii) $m \in C$ is *completely meet-irreducible* if for any $Y \subseteq C$, if $m = \bigwedge Y$ then $m = y$ for some $y \in Y$;
- (iii) $j \in C$ is *completely join-prime* if for any $X \subseteq C$, if $j \leq \bigvee X$ then $j \leq x$ for some $x \in X$;
- (iv) $m \in C$ is *completely meet-prime* if for any $Y \subseteq C$, if $m \geq \bigwedge Y$, then $m \geq y$ for some $y \in Y$.

A *perfect lattice* is a complete lattice in which the completely join-irreducible elements are join-dense (i.e., every element is a join of join-irreducibles), and the completely meet-irreducible elements are meet-dense (i.e., every element is a meet of meet-irreducibles). A *perfect distributive lattice* is a perfect lattice that is also completely distributive, i.e., arbitrary meets distribute over arbitrary joins and vice versa. In this case the completely join-irreducible (completely meet-irreducible) elements coincide with the completely join-prime (completely meet-prime) elements. (In general, a completely join-prime (completely meet-prime) element of a complete lattice is completely join-irreducible (completely meet-irreducible) but not vice versa.)

A bi-Heyting algebra is *perfect* if its lattice reduct is a perfect distributive lattice. It follows that in a perfect bi-Heyting algebra \mathbf{A} , for any $S \subseteq A$, we have that $\bigvee S \rightarrow a = \bigwedge_{s \in S} (s \rightarrow a)$, $a \rightarrow \bigwedge S = \bigwedge_{s \in S} (a \rightarrow s)$, $\bigvee S - a = \bigvee_{s \in S} (s - a)$ and $a - \bigwedge S = \bigvee_{s \in S} (a - s)$. A *perfect modal bi-Heyting algebra* is a modal bi-Heyting algebra the bi-Heyting reduct of which is a perfect bi-Heyting algebra, and moreover such that \Box and \Diamond preserve arbitrary meets and joins, respectively. The latter property allows us to add to any perfect modal bi-Heyting algebra the adjoint operations \blacklozenge and \blacksquare uniquely defined by the inequalities:

$$\Diamond a \leq b \iff a \leq \blacksquare b \quad \text{and} \quad a \leq \Box b \iff \blacklozenge a \leq b.$$

As usual we say that maps $f : \mathbf{A} \rightarrow \mathbf{B}$ and $g : \mathbf{B} \rightarrow \mathbf{A}$ form an *adjoint pair* if, for all $a \in A$ and $b \in B$, it holds that $f(a) \leq b$ iff $a \leq g(b)$. Here f is the *left adjoint* and g the *right adjoint*. It is well known that a map between complete lattices is a left (right) adjoint iff it is completely join-preserving (completely meet-preserving).

A map $f : \mathbf{A}^n \rightarrow \mathbf{A}$ is the *left residual in the i -th coordinate* of a map $g_i : \mathbf{A}^n \rightarrow \mathbf{A}$ if, for all $a_1, \dots, a_n, b \in A$, it holds that $f(a_1, \dots, a_n) \leq b$ iff $g_i(a_1, \dots, a_{i-1}, b, a_{i+1}, \dots, a_n) \leq a_i$. Here g_i is the *right residual of f in the i -th coordinate*. It is easy to check that, if \mathbf{A} is a complete lattice, then f has a right-residual (left-residual) in the i -th coordinate iff it is completely join-preserving (completely meet-preserving) in that coordinate. Recall that if $f : \mathbf{A}^n \rightarrow \mathbf{A}$ is completely join-preserving (completely meet-preserving) then it preserves all non-empty joins (meets) in each coordinate, but need not preserve empty joins (meets) in each coordinate. Thus $f : \mathbf{A}^n \rightarrow \mathbf{A}$ being an adjoint does not guarantee that it is a residual coordinatewise, nor vice versa.

The canonical extension. Let \mathbf{A} be a bounded distributive lattice with additional operations, in particular, \mathbf{A} could be a modal bi-Heyting algebra. The canonical extension of \mathbf{A} , defined by Gehrke and Harding [8] and denoted \mathbf{A}^δ , is a perfect bounded distributive lattice which, up to an isomorphism fixing \mathbf{A} , is the unique extension in which \mathbf{A} is dense and compact:

density every element of \mathbf{A}^δ can be written both as a join of meets and as a meet of joins of elements from \mathbf{A} .

compactness for all $S, T \subseteq A$, if $\bigvee S \leq \bigwedge T$ then $\bigvee S' \leq \bigwedge T'$ for some finite $S' \subseteq S$ and $T' \subseteq T$.

Additional operations on \mathbf{A} can be extended to \mathbf{A}^δ in a standard way. For more details the reader is referred to the appendix.

Languages and their interpretations. Let PROP, FVAR, PHVAR, NOM and CNOM be disjoint sets of propositional variables, fixed point variables, placeholder variables, nominals and co-nominals, respectively. Formulas in the basic language \mathcal{L} of modal bi-Heyting algebras are defined recursively by

$$\varphi ::= \perp \mid \top \mid p \mid X \mid \varphi \wedge \psi \mid \varphi \vee \psi \mid \varphi \rightarrow \psi \mid \varphi - \psi \mid \Diamond \varphi \mid \Box \varphi$$

where $p \in \text{PROP}$ and $X \in \text{FVAR}$. We identify the language with its set of formulas/terms. Formulas in the extended language \mathcal{L}^+ are defined by

$$\varphi ::= \perp \mid \top \mid p \mid X \mid \mathbf{j} \mid \mathbf{m} \mid \varphi \wedge \psi \mid \varphi \vee \psi \mid \varphi \rightarrow \psi \mid \varphi - \psi \mid \Diamond \varphi \mid \Box \varphi \mid \blacksquare \varphi \mid \blacklozenge \varphi$$

where $p \in \text{PROP}$, $X \in \text{FVAR}$, $\mathbf{j} \in \text{NOM}$ and $\mathbf{m} \in \text{CNOM}$. Placeholder variables from PHVAR, denoted x, y, z , will be used as generic variables which can take on the roles of propositional and fixed point variables. They will also be used to enhance the clarity of the exposition when dealing with substitution instances of formulas.

On perfect modal bi-Heyting algebras \blacksquare and \blacklozenge are interpreted as the right and left adjoints of \Diamond and \Box , respectively. Elements of NOM (CNOM) are interpreted as elements of $J^\infty(\mathbf{A}^\delta)$ ($M^\infty(\mathbf{A}^\delta)$). We now describe two extensions of \mathcal{L} obtained by adding fixed point operators. The distinction between the two extensions will become clear when we define their interpretations on distributive lattices with operators.

We define \mathcal{L}_1 to be the set of terms which extends \mathcal{L} by allowing terms $\mu x.t(x)$ and $\nu x.t(x)$ where $t \in \mathcal{L}_1$, $x \in \text{FVAR}$ and $t(x)$ is positive in x . The second extension is denoted \mathcal{L}_2 and extends \mathcal{L} by allowing construction of the terms $\mu_2 x.t(x)$ and $\nu_2 x.t(x)$ where $t \in \mathcal{L}_2$, $x \in \text{FVAR}$ and $t(x)$ is positive in x .

The terms of \mathcal{L} are interpreted as usual on modal bi-Heyting algebras. The additional terms of \mathcal{L}_1 and \mathcal{L}_2 are interpreted as follows: Suppose $t(x_1, x_2, \dots, x_n) \in \mathcal{L}_1$ and $a_1, \dots, a_{n-1} \in A$. Then

$$\mu x.t(x, a_1, \dots, a_{n-1}) := \bigwedge \{ a \in A \mid t(a, a_1, \dots, a_{n-1}) \leq a \}$$

if this meet exists, otherwise $\mu x.t(x, a_1, \dots, a_{n-1})$ is undefined. Similarly,

$$\nu x.t(x, a_1, \dots, a_{n-1}) := \bigvee \{ a \in A \mid a \leq t(a, a_1, \dots, a_{n-1}) \}$$

if this join exists, otherwise $\nu x.t(x, a_1, \dots, a_{n-1})$ is undefined. For each ordinal α we define $t^\alpha(\perp, a_2, \dots, a_n)$ as follows:

$$\begin{aligned} t^0(\perp, a_1, \dots, a_{n-1}) &= \perp, & t^{\alpha+1}(\perp, a_1, \dots, a_{n-1}) &= t(t^\alpha(\perp, a_1, \dots, a_{n-1}), a_1, \dots, a_{n-1}), \\ t^\lambda(\perp, a_1, \dots, a_{n-1}) &= \bigvee_{\alpha < \lambda} t^\alpha(\perp, a_1, \dots, a_{n-1}) && \text{for limit ordinals } \lambda; \\ t_0(\top, a_1, \dots, a_{n-1}) &= \top, & t_{\alpha+1}(\top, a_1, \dots, a_{n-1}) &= t(t_\alpha(\top, a_1, \dots, a_{n-1}), a_1, \dots, a_{n-1}), \\ t_\lambda(\top, a_1, \dots, a_{n-1}) &= \bigwedge_{\alpha < \lambda} t_\alpha(\top, a_1, \dots, a_{n-1}) && \text{for limit ordinals } \lambda. \end{aligned}$$

For $t(x_1, \dots, x_n) \in \mathcal{L}_2$ we then define

$$\mu_2 x.t(x, a_1, \dots, a_{n-1}) := \bigvee_{\alpha \geq 0} t^\alpha(\perp, a_1, \dots, a_{n-1}) \quad \text{and} \quad \nu_2 x.t(x, a_1, \dots, a_{n-1}) := \bigwedge_{\alpha \geq 0} t^\alpha(\top, a_1, \dots, a_{n-1})$$

if this join and this meet exist, and they are undefined otherwise.

A modal bi-Heyting algebra \mathbf{A} is said to be *of the first kind* (of the *second kind*) if $t^A(a_1, \dots, a_n)$ is defined for all $a_1, \dots, a_n \in \mathbf{A}$ and all $t \in \mathcal{L}_1$ ($t \in \mathcal{L}_2$). Henceforth we will refer to these algebras as *mu-algebras of the first kind* (of the *second kind*). When restricted to the Boolean case, our mu-algebras of the first kind are essentially the modal mu-algebras defined in [3, Definition 2.2] and [1, Definition 5.1].

Lemma 2.2. [1, Proposition 2.4] *If \mathbf{A} is a mu-algebra of the second kind, then \mathbf{A} is a mu-algebra of the first kind.*

Proof. Suppose $t(x, x_1, \dots, x_{n-1}) \in \mathcal{L}_2$ is positive in x and let $a_1, \dots, a_{n-1} \in A$. Further, let γ be the first ordinal such that $t^\beta(\perp, a_1, \dots, a_{n-1}) = t^\gamma(\perp, a_1, \dots, a_{n-1})$ for all $\beta > \gamma$ — such a γ exists since t is monotone in x . We will show that the meet $\mu x.t(x, a_1, \dots, a_{n-1}) = \bigwedge \{a \in A \mid t(a, a_1, \dots, a_{n-1}) \leq a\}$ exists by showing that $\mu_2 x.t(x, a_1, \dots, a_{n-1}) = \mu x.t(x, a_1, \dots, a_{n-1})$.

Suppose that $a \in A$ is a pre-fixed point of t , that is, $t(a, a_1, \dots, a_{n-1}) \leq a$. We will prove by transfinite induction that $t^\alpha(\perp, a_1, \dots, a_{n-1}) \leq a$ for all ordinals α .

- Base case: clearly $t^0(\perp, a_1, \dots, a_{n-1}) = \perp \leq a$.
- Suppose $t^\alpha(\perp, a_1, \dots, a_{n-1}) \leq a$. Then $t^{\alpha+1}(\perp, a_1, \dots, a_{n-1}) = t(t^\alpha(\perp, a_1, \dots, a_{n-1}), a_1, \dots, a_{n-1}) \leq t(a, a_1, \dots, a_{n-1}) \leq a$ since t is positive in x .
- Let λ be a limit ordinal with $t^\alpha(\perp, a_1, \dots, a_{n-1}) \leq a$ for all $\alpha < \lambda$. Then $t^\lambda(\perp, a_1, \dots, a_{n-1}) = \bigvee \{t^\alpha(\perp, a_1, \dots, a_{n-1}) \mid \alpha < \lambda\} \leq a$ since a is an upper bound for the set $\{t^\alpha(\perp, a_1, \dots, a_{n-1}) \mid \alpha < \lambda\}$.

Thus we have that $t^\gamma(\perp, a_1, \dots, a_{n-1}) \leq \bigwedge \{a \in A \mid t(a, a_1, \dots, a_{n-1}) \leq a\}$. Note also that $t^\gamma(\perp, a_1, \dots, a_{n-1})$ is a pre-fixed point as $t(t^\gamma(\perp, a_1, \dots, a_{n-1}), a_1, \dots, a_{n-1}) = t^{\gamma+1}(\perp, a_1, \dots, a_{n-1}) = t^\gamma(\perp, a_1, \dots, a_{n-1})$. Thus $t^\gamma(\perp, a_1, \dots, a_{n-1}) \in \{a \in A \mid t(a, a_1, \dots, a_{n-1}) \leq a\}$ and so we have the desired equality. \square

The importance of Lemma 2.2 is that if we are interpreting formulas/terms on a mu-algebra of the second kind the interpretations of terms with the two different fixed point binders will agree. That is, $\mu X.\varphi(X) = \mu_2 X.\varphi(X)$ and $\nu X.\psi(X) = \nu_2 X.\psi(X)$.

The final sets of terms, \mathcal{L}_* (respectively, \mathcal{L}_*^+), are obtained as an extension of \mathcal{L} (respectively, \mathcal{L}^+) by allowing $\mu^* x.t(x)$ and $\nu^* x.t(x)$ whenever $t \in \mathcal{L}_*$ (respectively, $t \in \mathcal{L}_*^+$) and positive in x . Terms in \mathcal{L}_* and \mathcal{L}_*^+ are only interpreted in the canonical extensions \mathbf{A}^δ of modal bi-Heyting algebras \mathbf{A} . If $t(x_1, x_2, \dots, x_n) \in \mathcal{L}_* \cup \mathcal{L}_*^+$ and $a_1, \dots, a_{n-1} \in A^\delta$, then $\mu^* x_1.t(x_1, a_1, \dots, a_{n-1}) := \bigwedge \{a \in A \mid t(a, a_1, \dots, a_{n-1}) \leq a\}$ and $\nu^* x_1.t(x_1, a_2, \dots, a_n) := \bigvee \{a \in A \mid a \leq t(a, a_1, \dots, a_{n-1})\}$. As the canonical extension \mathbf{A}^δ is a complete lattice, the interpretation of $\mu^* x.t(x)$ or $\nu^* x.t(x)$ is always defined. Given a term $\varphi \in \mathcal{L}_1^+$ we write φ^* for the \mathcal{L}_*^+ term obtained from φ by replacing all occurrences of μ and ν with μ^* and ν^* , respectively.

An \mathcal{L}^+ formula is *pure* if it contains no ordinary (propositional) variables but only, possibly, nominals and co-nominals. A formula of \mathcal{L}_1 (respectively \mathcal{L}_1^+ , \mathcal{L}_2 , \mathcal{L}_2^+ , \mathcal{L}_* , \mathcal{L}_*^+) is an \mathcal{L}_1 -sentence if it contains no free fixed point variables (and similarly for \mathcal{L}_1^+ , \mathcal{L}_2 , \mathcal{L}_2^+ , \mathcal{L}_* , \mathcal{L}_*^+).

The reason for using μ^* and ν^* in formulas/terms is so that the interpretation does not change when moving between \mathbf{A} and \mathbf{A}^δ .

Quasi-inequalities, assignments, validity. An assignment on \mathbf{A} sends propositional variables to elements of \mathbf{A} and is extended to formulas of \mathcal{L}_1 , \mathcal{L}_2 and \mathcal{L}_* in the usual way, where these are defined. An assignment on \mathbf{A}^δ sends propositional variables to elements of \mathbf{A}^δ , nominals into $J^\infty(\mathbf{A}^\delta)$ and co-nominals into $M^\infty(\mathbf{A}^\delta)$ and extends to all formulas of \mathcal{L}_*^+ , \mathcal{L}_1^+ and \mathcal{L}_2^+ . An *admissible assignment* on \mathbf{A}^δ is an assignment which takes all propositional variables to elements of \mathbf{A} . An \mathcal{L}^+ -inequality $\alpha \leq \beta$ is *admissibly valid* on \mathbf{A}^δ , denoted $\mathbf{A}^\delta \models_{\mathbf{A}} \alpha \leq \beta$, if it holds under all admissible assignments.

A *quasi-inequality* of \mathcal{L}_1 (resp., \mathcal{L}_1^+ , \mathcal{L}_2 , \mathcal{L}_2^+ , \mathcal{L}_* , \mathcal{L}_*^+) is an expression of the form $\varphi_1 \leq \psi_1 \& \dots \& \varphi_n \leq \psi_n \Rightarrow \varphi \leq \psi$ where the φ_i , ψ_i , φ and ψ are formulas of \mathcal{L}_1 (resp., \mathcal{L}_1^+ , \mathcal{L}_2 , \mathcal{L}_2^+ , \mathcal{L}_* , \mathcal{L}_*^+). A quasi-inequality $\varphi_1 \leq \psi_1 \& \dots \& \varphi_n \leq \psi_n \Rightarrow \varphi \leq \psi$ is satisfied under an assignment V in an algebra \mathbf{A} of the appropriate sort, written $\mathbf{A}, V \models \varphi_1 \leq \psi_1 \& \dots \& \varphi_n \leq \psi_n \Rightarrow \varphi \leq \psi$ if $\mathbf{A}, V \not\models \varphi_i \leq \psi_i$ for some $1 \leq i \leq n$ or $\mathbf{A}, v \models \varphi \leq \psi$. A quasi-inequality is (admissibly) valid in an algebra if it is satisfied by every (admissible) assignment.

Signed generation trees To any formula/term in \mathcal{L}_1^+ and \mathcal{L}_*^+ we assign two *signed generation trees*. That is, for $\varphi \in \mathcal{L}$ we consider two trees $+\varphi$ and $-\varphi$. The generation tree is constructed as usual, beginning at the root with the main connective and then branching out into n -nodes at each n -ary connective. Each leaf is either a propositional variable, a fixed point variable, or a constant. Each node is signed as follows:

- the root node of $+\varphi$ is signed $+$ and the root node of $-\varphi$ is signed $-$;
- if a node is \vee , \wedge , \diamond , \square , \blacklozenge , or \blacksquare assign the same sign to its successor nodes;
- if a node is \rightarrow , assign the opposite sign to its left successor, and the same sign to its right successor;
- if a node is $-$, assign the same sign to its left successor, and the opposite sign to its right successor;
- if a node is $\mu x.\varphi(x)$, $\mu^* x.\varphi(x)$, $\nu x.\varphi(x)$ or $\nu^* x.\varphi(x)$ (with every free occurrence of x in the positive generation tree of φ labelled positively) then assign the same sign to the successor node.

A node in a signed generation tree is said to be *positive* if it is signed “ $+$ ” and *negative* if it is signed “ $-$ ”. Examples of signed generation trees can be found in Figure 2 and Figure 3.

Order types We will often be using formulas in n variables and hence use \bar{x} to denote n -tuple of variables. An *order-type* over $n \in \mathbb{N}$ is an n -tuple $\epsilon \in \{1, \partial\}^n$. Given an order-type ϵ , its opposite order type, denoted ϵ^∂ , is given by $\epsilon_i^\partial = 1 \Leftrightarrow \epsilon_i = \partial$ for $1 \leq i \leq n$. We will also use the symbol τ to denote an order type over n . When we define the Approximation Rules in Section 5 it will be useful to write \top^1 and \top^∂ for \top and \perp respectively. Similarly we will write \perp^1 and \perp^∂ for \perp and \top respectively.

For both order types and tuples of variables, we will use the symbol \oplus to denote concatenation.

Join- and meet-irreducible elements in products Let \mathbf{A} be a perfect lattice and τ and order type over n . It is not difficult to prove that in the product \mathbf{A}^n every join-irreducible element \bar{j} is an n -tuple such that $(\bar{j})_k \in J^\infty(\mathbf{A})$ for some $1 \leq k \leq n$ and $(\bar{j})_i = \perp$ for all $i \neq k$. Dually, every completely meet-irreducible element in the product \mathbf{A}^n is an n -tuple \bar{m} such that $(\bar{m})_k \in M^\infty(\mathbf{A})$ for some $1 \leq k \leq n$ and $(\bar{m})_i = \top$ for all $i \neq k$.

This characterization of $J^\infty(\mathbf{A}^n)$ and $M^\infty(\mathbf{A}^n)$ can be generalized to the case of $J^\infty(\mathbf{A}^\tau)$ and $M^\infty(\mathbf{A}^\tau)$: every join-irreducible element $\bar{j} \in J^\infty(\mathbf{A}^\tau)$ is an n -tuple such that, for some $1 \leq k \leq n$, $(\bar{j})_k \in J^\infty(\mathbf{A})$ if $\tau_k = 1$ and $(\bar{j})_k \in M^\infty(\mathbf{A})$ if $\tau_k = \partial$ while, for $i \neq k$, $(\bar{j})_i = \perp$ if $\tau_i = 1$ and $(\bar{j})_i = \top$ if $\tau_i = \partial$. Order-dually, every meet-irreducible element $\bar{m} \in M^\infty(\mathbf{A}^\tau)$ is an n -tuple such that, for some $1 \leq k \leq n$, $(\bar{m})_k \in M^\infty(\mathbf{A})$ if $\tau_k = 1$ and $(\bar{m})_k \in J^\infty(\mathbf{A})$ if $\tau_k = \partial$ while, for $i \neq k$, $(\bar{m})_i = \top$ if $\tau_i = 1$ and $(\bar{m})_i = \perp$ if $\tau_i = \partial$.

Let \mathbf{A} be a perfect lattice and let τ be an order type over n . The notation \mathbf{j}^{τ_i} will denote a nominal if $\tau_i = 1$ and a co-nominal if $\tau_i = \partial$. Dually, \mathbf{n}^{τ_i} denotes a co-nominal if $\tau_i = 1$ and a nominal if $\tau_i = \partial$.

Accordingly, we will use the notation $\bar{\mathbf{j}}_i^\tau$ to denote an n -tuple in which the i -th component is \mathbf{j}^{τ_i} and, for all $k \neq i$, the k -th component is \perp^{τ_k} . Accordingly, such tuples range over the subset of $J^\infty(\mathbf{A}^\tau)$ in which the i -th

$$\begin{array}{ccc}
\mathbf{A} \models \alpha \leq \beta & & \mathbf{A}^\delta \models \alpha \leq \beta \\
\Downarrow & & \Downarrow \\
\mathbf{A}^\delta \models_{\mathbf{A}} \alpha \leq \beta & & \\
\Downarrow & & \\
\mathbf{A}^\delta \models_{\mathbf{A}} \text{pure}(\alpha \leq \beta) & \iff & \mathbf{A}^\delta \models \text{pure}(\alpha \leq \beta)
\end{array}$$

Figure 1: The U-shaped argument for canonicity of inequalities interpreted on a lattice-based algebra \mathbf{A} .

component comes from $J^\infty(\mathbf{A}^{\tau_i})$ and all other components are \perp^{τ_k} , $k \neq i$. Dually, $\bar{\mathbf{n}}_i^\tau$ denotes an n -tuple in which the i -th component is \mathbf{n}^{τ_i} and, for all $k \neq i$, the k -th component is \top^{τ_k} .

3 Canonicity results

In this section we highlight and explain the main results of this paper and describe our methodology in broad strokes. Our approach goes via a “U-shaped argument” (see Figure 1), a generic version of which we now outline.

Going down the left-hand arm of the diagram, the first bi-implication is given by the fact that validity in \mathbf{A} coincides with admissible validity in \mathbf{A}^δ , modulo certain provisions regarding the fixed point binders. In the richer setting of \mathbf{A}^δ we can now interpret the extended language \mathcal{L}_*^+ where equivalences involving the adjoints \blacklozenge and \blacksquare as well as nominals and co-nominals are available. The aim is now to transform the inequality into a set of pure (quasi-)inequalities, denoted $\text{pure}(\alpha \leq \beta)$ in Figure 1. This is done by means of an algorithm, μ^* -ALBA, based on a calculus of rewrite rules which are presented in Section 5. The fact that admissible and ordinary validity coincide for pure inequalities is the linchpin for the transition from validity in \mathbf{A} (simulated as admissible validity in \mathbf{A}^δ) to validity in \mathbf{A}^δ . This justifies the bi-implication forming the base of the “U”.

We progress up the right-hand arm of the ‘U’ by reversing the rewrite rules applied when coming down the left-hand side. The equivalences are justified by the fact that these rules preserve validity on perfect algebras. (The last statement needs to be qualified somewhat to accommodate fixed point binders, as will be specified in Section 9.) We observe that the equivalences on the right-hand arm of Figure 1 justifies the first-order frame definability of $\alpha \leq \beta$, which is the topic of [5]. Indeed, the fact that $\alpha \leq \beta$ is equivalent on perfect algebras (or, dually, relational structures) to a set of pure quasi-inequalities guarantees first-order definability since the absence of propositional variables in $\text{pure}(\alpha \leq \beta)$ means that all quantification ranges over first-order definable subsets of the dual relational structure of \mathbf{A}^δ .

As we now explain, the notion of canonicity of a formula in the presence of fixed points admits some variation. If φ is a formula without fixed point binders, then the term function $\varphi^{\mathbf{A}^\delta}$ extends the term function $\varphi^{\mathbf{A}}$, i.e., they agree on arguments from \mathbf{A} . This is something which is usually of crucial importance in proving that an equation is canonical. As soon as we add fixed point binders this extension property fails. Indeed, $(\varphi(X))^{\mathbf{A}^\delta}$ can have more pre-fixed points in \mathbf{A}^δ than $(\varphi(X))^{\mathbf{A}}$ has in \mathbf{A} , and so $(\mu X.\varphi(X))^{\mathbf{A}^\delta}$ would generally be smaller than $(\mu X.\varphi(X))^{\mathbf{A}}$. This phenomenon creates significant obstacles for standard canonicity arguments.

One possible remedy to regain the extension property, is to insist that only pre-fixed points from the smaller algebra \mathbf{A} be considered in calculating $(\mu X.\varphi(X))^{\mathbf{A}^\delta}$, i.e., rather calculating $(\mu^* X.\varphi(X))^{\mathbf{A}^\delta}$. This leads to the first notion of canonicity that we will examine. This is the notion described by Bezhanishvili and Hodkinson [3], and used by them to obtain completeness results for some axiomatic extensions of the basic (classical) mu-calculus with respect to certain types of general frames. We will refer to this as *tame canonicity*. In tame canonicity, all fixed point binders μ and ν in inequalities are replaced by the binders μ^* and ν^* and interpreted as such in the canonical extension. Proving that an \mathcal{L}_1 -inequality is *tame canonical* is proving that $\mathbf{A} \models \varphi \leq \psi$ if and only if $\mathbf{A}^\delta \models \varphi^* \leq \psi^*$ where \mathbf{A} is mu-algebra of the first kind. (Further explanation regarding the use of μ^* and ν^* is

given at the end of Section 7.)

Using our algorithmic approach we are able to prove tame canonicity for what we call *tame inductive* mu-inequalities. This is a smaller class than that for which correspondence results were shown in [5] but extends the Sahlqvist mu-inequalities of [3] when projected to the classical case.

The second type of canonicity which we investigate is essentially the usual notion of canonicity and thus we simply refer to this as *canonicity*. That is, for \mathbf{A} a mu-algebra of the second kind, if $\mathbf{A} \models \varphi \leq \psi$, then $\mathbf{A}^\delta \models \varphi \leq \psi$. (Note the additional assumption that \mathbf{A} is of the second kind.) Our method of proof is as follows: we first show that $\mathbf{A} \models \varphi \leq \psi$ if and only if $\mathbf{A}^\delta \models_{\mathbf{A}} \varphi^* \leq \psi^*$ and then that this implies $\mathbf{A}^\delta \models \varphi \leq \psi$. However, we are not able to show the converse, i.e., that $\mathbf{A}^\delta \models \varphi \leq \psi$ implies $\mathbf{A} \models \varphi \leq \psi$.

We will call the syntactically specified class of formulas for which we can prove canonicity the *restricted inductive mu-inequalities*. This class is again a restriction of the class for which correspondence results were shown in [5], but it is a generalization of the inductive inequalities from [6].

From a logical perspective, one of the main motivations for proving canonicity results is to obtain relational completeness results for axiomatic extension of logics — if we have a general algebraic completeness result any non-theorem is refuted on an algebra of the logic and, if the axioms are preserved under canonical extensions, we can transfer this refutation to a relational structure obtained as the dual of the perfect algebra which is the canonical extension. Unfortunately this does not work in the setting of the mu-calculus, since generally it will not be true that $\mathbf{A}^\delta \models \varphi \leq \psi$ implies that $\mathbf{A} \models \varphi \leq \psi$ or, contrapositively, that $\mathbf{A} \not\models \varphi \leq \psi$ implies $\mathbf{A}^\delta \not\models \varphi \leq \psi$. Hence refutations will not necessarily be preserved when taking canonical extensions.

The proof of the tame canonicity of the tame inductive mu-inequalities and the proof of the canonicity of the restricted inductive mu-inequalities is concluded in Section 10. The proof is a two-step process, beginning in Section 9. There it is shown that whenever a *tame* run of μ^* -ALBA succeeds on a mu-inequality $\varphi \leq \psi$, we have that $\varphi \leq \psi$ is *tame canonical*. It is also shown that whenever a *proper* run succeeds on a mu-inequality $\alpha \leq \beta$, then $\alpha \leq \beta$ will be *canonical*. Section 10 then completes the overall result by showing that for every tame inductive mu-inequality (respectively, a restricted inductive mu-inequality), there exists a tame (respectively, proper) run of μ^* -ALBA which succeeds on that inequality.

4 Syntactic classes

In this section we introduce two new syntactically defined classes of mu-inequalities. Both will be subclasses of the recursive mu-inequalities introduced in [5], and therefore the members of both classes will all have first-order frame correspondents.

The reason we concern ourselves with *inequalities* is because certain fragments (for example, distributive modal logic [9]) do not have an implication in the language. In the absence of an implication, modal sequents of the form $\varphi \Rightarrow \psi$ are used to capture the logic. When using the algebraic semantics, the interpretation of inequality $\varphi \leq \psi$ is used as the interpretation of the sequent $\varphi \Rightarrow \psi$.

For any \mathcal{L}_1 -sentence $\varphi(p_1, \dots, p_n)$, any order-type ϵ over n , and any $1 \leq i \leq n$, an ϵ -critical node in a signed generation tree of φ is a (leaf) node $+p_i$ with $\epsilon_i = 1$, or $-p_i$ with $\epsilon_i = \partial$. An ϵ -critical branch in the tree is a branch terminating in an ϵ -critical node. The intuition, which will be built upon later, is that variable occurrences corresponding to ϵ -critical nodes are *to be solved for, according to ϵ* .

In the signed generation tree of a \mathcal{L}_1 -sentence $\varphi(p_1, \dots, p_n)$ a *live branch* is a branch ending in a (signed) propositional variable. In particular, all critical branches are live. It follows that a branch is not live iff it ends in a propositional constant (\top or \perp) or in a fixed point variable.

For every \mathcal{L}_1 -sentence $\varphi(p_1, \dots, p_n)$, and every order-type ϵ , we say that $+\varphi$ (resp. $-\varphi$) *agrees with ϵ* , and write $\epsilon(+\varphi)$ (resp. $\epsilon(-\varphi)$), if every leaf node in the signed generation tree of $+\varphi$ (resp. $-\varphi$) which is labelled with a propositional variable is ϵ -critical. In other words, $\epsilon(+\varphi)$ (resp. $\epsilon(-\varphi)$) means that all propositional variable occurrences corresponding to leaves of $+\varphi$ (resp. $-\varphi$) are to be solved for according to ϵ . We will also make use of the *sub-tree relation* $\gamma < \varphi$, which extends to signed generation trees, and we will write $\epsilon(\gamma) < *\varphi$ to indicate that γ , regarded as a sub- (signed generation) tree of $*\varphi$, agrees with ϵ .

Outer Skeleton (P_3)	Inner Skeleton (P_2)	PIA (P_1)
Δ -adjoints	Binders	Binders
$+$ \vee \wedge	$+$ μ	$+$ ν
$-$ \wedge \vee	$-$ ν	$-$ μ
SLR	SLA	SRA
$+$ \diamond \triangleleft \circ $-$	$+$ \diamond \triangleleft \vee	$+$ \square \triangleright \wedge
$-$ \square \triangleright \star \rightarrow	$-$ \square \triangleright \wedge	$-$ \diamond \triangleleft \vee
	SLR	SRR
	$+$ \wedge \circ $-$	$+$ \vee \star \rightarrow
	$-$ \vee \star \rightarrow	$-$ \wedge \circ $-$

Table 1: Skeleton and PIA nodes.

While reading the following definition, the reader might find it useful to refer to Example 4.3 for an illustration of the concepts being introduced.

Definition 4.1. Nodes in signed generation trees will be called *skeleton nodes* and *PIA nodes* and further classified as Δ -adjoint, SLR, Binders, SLA, SRA or SRR, according to the specification given in table 1.¹ The acronym *PIA* stands for “positive implies atomic”.

Definition 4.2. Let $\varphi(p_1, \dots, p_n)$ be a formula in the propositional variables p_1, \dots, p_n , and let ϵ be an order type on $\{1, \dots, n\}$.

A branch in a signed generation tree $*\varphi$, for $*$ $\in \{+, -\}$, ending in a propositional variable is an ϵ -good branch if, apart from the leaf, it is the concatenation of three paths P_1 , P_2 , and P_3 , each of which may possibly be of length 0, such that P_1 is a path from the leaf consisting only of PIA-nodes, P_2 consists only of inner skeleton-nodes, and P_3 consists only of outer skeleton-nodes and, moreover, it satisfies conditions (GB1), (GB2) and (GB3), below.

(GB1) The formula corresponding to the uppermost node on P_1 is a mu-sentence.

(GB2) For every SRR-node in P_1 of the form $\gamma \odot \beta$ or $\beta \odot \gamma$, where β is the side where the branch lies, γ is a mu-sentence and $\epsilon^\partial(\gamma) < *\varphi$ (i.e., γ contains no variable occurrences to be solved for — see above).

Unravelling the condition $\epsilon^\partial(\gamma) < *\varphi$ specifically to the \mathcal{L}_1 -signature (expanded with \circ and \star), we obtain:

- a) if $\gamma \odot \beta$ is $+(\gamma \star \beta)$, $+(\gamma \vee \beta)$, $+(\beta \rightarrow \gamma)$, or $-(\beta - \gamma)$, then $\epsilon^\partial(+\gamma)$;
- b) if $\gamma \odot \beta$ is $+(\gamma \rightarrow \beta)$, $-(\gamma \wedge \beta)$, $-(\gamma \circ \beta)$, or $-(\gamma - \beta)$, then $\epsilon^\partial(-\gamma)$ (equivalently, $\epsilon(+\gamma)$).

(GB3) For every SLR-node in P_2 of the form $\gamma \odot \beta$ or $\beta \odot \gamma$, where β is the side where the branch lies, γ is a mu-sentence and $\epsilon^\partial(\gamma) < *\varphi$ (see above for this notation).

Unravelling the condition $\epsilon^\partial(\gamma) < *\varphi$ specifically to the \mathcal{L}_1 -signature (expanded with \circ and \star), we obtain:

- a) if $\gamma \odot \beta$ is $-(\gamma \star \beta)$, $-(\gamma \vee \beta)$, $-(\beta \rightarrow \gamma)$, or $+(\beta - \gamma)$, then $\epsilon^\partial(-\gamma)$ (equivalently, $\epsilon(+\gamma)$);
- b) if $\gamma \odot \beta$ is $-(\gamma \rightarrow \beta)$, $+(\gamma \wedge \beta)$, $+(\gamma \circ \beta)$, or $+(\gamma - \beta)$, then $\epsilon^\partial(+\gamma)$.

As promised, the example below illustrates the concept of an ϵ -good branch.

Example 4.3. The generation trees of $+\diamond\mu X.(\diamond X \vee \square(\square\diamond q \vee p))$ and $-\nu Y.([\square((q \rightarrow \perp) \wedge (p \rightarrow \perp)) \rightarrow \perp] \wedge \square Y)$ are given in Figure 2. Taking ϵ to be the order type with $\epsilon_p = 1$ and $\epsilon_q = \partial$ there are two ϵ -critical branches,

¹The interpretations of the binary connectives *fusion* \circ and *fission* \star included in table 1 preserve, respectively, joins and meets coordinatewise. The unary connectives \triangleleft and \triangleright are, respectively meet- and join-reversing and were studied in the context of Distributive Modal Logic by Gehrke, Nagahashi and Venema [9], see also [6]. Although \circ , \star , \triangleleft and \triangleright are not part of our language, their inclusion here illustrates the fact that the classification and subsequent definitions are based solely on the order-theoretic behaviours of the interpretations of connectives and can hence be easily ported to other languages.




Figure 2: The generation trees of $+\Diamond\mu X.(\Diamond X \vee \Box(\Diamond q \vee p))$ and $-\nu Y.([\Box((q \rightarrow \perp) \wedge (p \rightarrow \perp)) \rightarrow \perp] \wedge \Box Y)$ with ϵ -good branches indicated for an order type ϵ with $\epsilon_p = 1$ and $\epsilon_q = \partial$.

namely the one ending in $+p$ and the one ending in $-q$. These are both ϵ -good. Indeed, they can correctly be split into P_1 , P_2 and P_3 paths as indicated in the figure. Let us verify that the branch ending in $+p$ moreover satisfies (GB1), (GB2) and (GB3): the formula $\Box(\Box\Diamond q \vee p)$ corresponding to the uppermost P_1 node is a mu-sentence as it contains no fixed point variables, so (GB1) holds. The only SRR-node on P_1 is $+\vee$ and we must check that it satisfies (GB2): here the role of γ is played by $\Box\Diamond q$ which, firstly, is a mu-sentence; secondly, the only occurring propositional variable is q which occurs positively while $\epsilon_q = \partial$, hence $\epsilon^\partial(\Box\Diamond q) < +\Diamond\mu X.(\Diamond X \vee \Box(\Box\Diamond q \vee p))$. There are no SLR nodes in P_2 , so (GB3) holds vacuously.

Turning our attention to the branch ending in $-q$: the formula $\Box((q \rightarrow \perp) \wedge (p \rightarrow \perp))$ corresponding to the uppermost P_1 node is a mu-sentence, and hence (GB1) is satisfied. There is one SLR-node, namely $-\rightarrow$, which verifies (GB3) as the formula playing the role of γ is the constant sentence \perp . Similarly, there is one SRR-node, $+\rightarrow$, which satisfies (GB2) as, in this instance, formula playing the role of γ is again the constant sentence \perp .

Our main interest is in ϵ -good branches satisfying some of the additional properties in the following definition.

Definition 4.4. Let $\varphi(p_1, \dots, p_n)$ be a formula in the propositional variables p_1, \dots, p_n , let ϵ be an order type on $\{1, \dots, n\}$ and $<_\Omega$ a strict partial order on the variables p_1, \dots, p_n . An ϵ -good branch may satisfy one or more of the following properties:

(NB-PIA) P_1 contains no fixed point binders.

(NL) For every SLR-node in P_2 of the form $\gamma \odot \beta$ or $\beta \odot \gamma$, where β is the side where the branch lies, the signed generation tree of γ contains no live branches. Unravelling this condition specifically to the \mathcal{L} -signature, we obtain:

- a) if $\gamma \odot \beta$ is $-(\gamma \star \beta)$, $-(\gamma \vee \beta)$, $-(\beta \rightarrow \gamma)$, or $+(\beta - \gamma)$, then $-\gamma$ contains no live branches;
- b) if $\gamma \odot \beta$ is $-(\gamma \rightarrow \beta)$, $+(\gamma \wedge \beta)$, $+(\gamma \circ \beta)$, or $+(\gamma - \beta)$, then $+\gamma$ contains no live branches.

(Ω -CONF) For every SRR-node in P_1 of the form $\gamma \odot \beta$ or $\beta \odot \gamma$, where β is the side where the branch lies: $p_j <_\Omega p_i$ for every p_j occurring in γ , where p_i is the propositional variable labelling the leaf of the branch.

Remark 4.5. 1. Since there is overlap between the signed connectives in the columns of Table 1, the borders between the P_1 , P_2 and P_3 parts of a good branch need not be uniquely determined. For strategic reasons, to be discussed below, it is preferable make P_3 as long as possible at the cost of P_2 and, in turn, make P_2 as long as possible at the cost of P_1 .

2. The abbreviations SLR, SLA, SRA and SRR stand for syntactically left residual, left adjoint, right adjoint and right residual, respectively. Nodes are thus classified according to the order-theoretic properties of their interpretations.

Definition 4.6. Given an order type ϵ and a strict partial order $<_{\Omega}$ on the variables p_1, \dots, p_n , the signed generation tree $*\varphi$, $*$ $\in \{-, +\}$, of a term $\varphi(p_1, \dots, p_n)$ is called

1. ϵ -recursive if every ϵ -critical branch is ϵ -good.
2. (Ω, ϵ) -inductive if it is ϵ -recursive and every ϵ -critical branch satisfies $(\Omega\text{-CONF})$.
3. *restricted* (Ω, ϵ) -inductive if it is (Ω, ϵ) -inductive and
 - (a) every ϵ -critical branch satisfies (NB-PIA) and (NL),
 - (b) every occurrence of a binder is on an ϵ -critical branch.
4. *tame* (Ω, ϵ) -inductive if it is (Ω, ϵ) -inductive and
 - (a) $\Omega = \emptyset$,
 - (b) no binder occurs on any ϵ -critical branch,
 - (c) the only nodes involving binders which are allowed to occur are $+\nu$ and $-\mu$.

An inequality $\varphi \leq \psi$ is ϵ -recursive (resp., (Ω, ϵ) -inductive, *restricted* (Ω, ϵ) -inductive, *tame* (Ω, ϵ) -inductive) if $+\varphi$ and $-\psi$ are both ϵ -recursive (resp., (Ω, ϵ) -inductive, *restricted* (Ω, ϵ) -inductive, *tame* (Ω, ϵ) -inductive).

An inequality $\varphi \leq \psi$ is recursive (resp., inductive, *restricted inductive*, *tame inductive*) if $\varphi \leq \psi$ is ϵ -recursive (resp., (Ω, ϵ) -inductive, *restricted* (Ω, ϵ) -inductive, *tame* (Ω, ϵ) -inductive) for some strict partial order Ω and order type ϵ .

The corresponding classes of inequalities will be referred to as the *recursive* (resp., *inductive*, *restricted inductive*, *tame inductive*) *mu-inequalities*, or the *recursive* (resp., *inductive*, *restricted inductive*, *tame inductive*) *mu-formulas*, if the inequality signs have been replaced with implications.

Remark 4.7. If we were to interpret \mathcal{L}_1 formulas classically, the tame inductive formulas would include the *Sahlqvist fixed point formulas* of [3, Definition 5.1]. In this setting the tame inductive formulas are, for example, slightly more liberal in terms of what they allow to play the role of a ‘boxed atom’ (see Example 4.10, below).

Example 4.8. Consider the inequality $\diamond\mu X.(\diamond X \vee \Box(\Box\diamond q \vee p)) \leq \nu Y.([\Box((q \rightarrow \perp) \wedge (p \rightarrow \perp)) \rightarrow \perp] \wedge \Box Y)$. The generation trees for $+\diamond\mu X.(\diamond X \vee \Box(\Box\diamond q \vee p))$ and $-\nu Y.([\Box((q \rightarrow \perp) \wedge (p \rightarrow \perp)) \rightarrow \perp] \wedge \Box Y)$ were given in Figure 2.

From the generation trees we can now see that this is a *restricted* (Ω, ϵ) -inductive inequality for ϵ such that $\epsilon_p = 1$ and $\epsilon_q = \partial$ and Ω such that $q <_{\Omega} p$. Indeed, given this ϵ , there are two ϵ -critical branches: the one ending in $+p$ and the one ending in $-q$. In Example 4.3 we verified that these branches are ϵ -good. It is straightforward to check that the branch ending in $+p$ also satisfies (NB-PIA) and (NL). We need to check that the SRR node $+\vee$ satisfies $(\Omega\text{-CONF})$: in this case $\Box\diamond q$ plays the role of γ , and indeed $\epsilon^{\partial}(\Box\diamond q)$ and $q <_{\Omega} p$.

Now considering the branch ending in $-q$: it is easy to check that it satisfies (NB-PIA) and (NL). To check that it satisfies $(\Omega\text{-CONF})$ we need to consider the SRR-node $+\rightarrow$, but here the role of γ is played by $+\perp$, so clearly $\epsilon^{\partial}(\gamma)$.

The inequality is *not* (Ω', ϵ') -inductive for any other order type ϵ' or dependency order Ω' . Indeed, we cannot have $\epsilon'_q = 1$, as the branch ending in $+q$ cannot be correctly divided in P_1 , P_2 and P_3 parts due the presence of $+\diamond$ in the scope of $+\Box$, and is therefore not good. Also we cannot have $\epsilon'_p = \epsilon'_q$, since in that case there would be a binder in one of the trees the occurrence of which is not on a critical branch, violating Definition 4.6.3(b). So since ϵ is the only possible order type, the configuration $+\Box(\Box\diamond q \vee p)$ and $(\Omega\text{-CONF})$ dictate that Ω is the only possible dependency order.





Figure 3: The generation trees of $+(\Diamond(\Box\perp \vee p) \wedge \Box q)$ and $-\mu Y.(\Diamond(p \wedge q) \wedge \Box Y)$ with ϵ -good branches indicated for an order type ϵ with $\epsilon_p = 1 = \epsilon_q$.

Example 4.9. Consider the formula $(\mu X.(p \vee \Diamond X) \wedge \mu X.(q \vee \Diamond X)) \rightarrow \mu X.((p \wedge \mu Y.(q \vee \Diamond Y)) \vee \Diamond X)$. This is valid at a point w in a frame iff all point reachable from w are reachable from one another. The corresponding inequality $(\mu X.(p \vee \Diamond X) \wedge \mu X.(q \vee \Diamond X)) \leq \mu X.((p \wedge \mu Y.(q \vee \Diamond Y)) \vee \Diamond X)$ is (Ω, ϵ) -inductive with $\epsilon_p = 1 = \epsilon_q$ (and for no other order type) and any strict partial order Ω . Note that it is not restricted inductive nor tame inductive, since binders occur on both critical and non-critical branches for every possible choice of ϵ .

Example 4.10. The inequality $\Diamond(\Box\perp \vee p) \wedge \Box q \leq \mu Y.(\Diamond(p \wedge q) \wedge \Box Y)$ is tame (Ω, ϵ) -inductive with $\epsilon_p = 1 = \epsilon_q$ (and $\Omega = \emptyset$), as can be seen from the generation trees in Figure 3. Moreover, the corresponding implication it is *not* a Sahlqvist fixed point formula in the sense of [3, Definition 5.1], as that definition (and the classical definition of Sahlqvist formulas) would not allow the configuration $\Box(\Box\perp \vee p)$ where p is a variable to be solved for. It is not tame (Ω, ϵ') -inductive for any ϵ' with $\epsilon'_p = \partial$ or $\epsilon'_q = \partial$, as in such cases the binder μ would appear on an ϵ -critical branch.

Finally note that this inequality is not restricted inductive. Indeed, since the definition of restricted (Ω, ϵ) -inductive inequalities requires all binders to occur on critical branches, the $-\mu$ would have to be on a critical branch, but then it would have to be on the P_1 part of that branch (according to the definition of good branches), yet (NB-PIA) prohibits this for restricted (Ω, ϵ) -inductive inequalities.

The relationships between the different classes of inequalities or, equivalently, formulas, under consideration are illustrated in the diagram below. The inclusions follow from the definitions. Each of the regions is non-empty: that regions A , B and F are non-empty follows from Examples 4.9, 4.8 and 4.10, respectively. Region C contains the inductive formulas which are not Sahlqvist and which contain no fixed point binders, e.g., the Frege axiom $(p \rightarrow (q \rightarrow r)) \rightarrow ((p \rightarrow q) \rightarrow (p \rightarrow r))$ (see [6, Examples 3.16 and 7.5]). Region D contains the binder free (intuitionistic) Sahlqvist formulas, e.g., $p \rightarrow \Box\Diamond p$, while the Sahlqvist μ -formulas which do contain binders are found in region E , e.g., $\Diamond p \leq \Box\mu X.(p \vee \Diamond X)$ (cf. [3, Example 6.4]).



5 The algorithm μ^* -ALBA

The algorithm μ -ALBA was introduced in [5, Section 1]. Here we present a variant called μ^* -ALBA, and which is a restricted version of μ -ALBA. The goal of our algorithm is to *eliminate propositional variables* from inequalities, while maintaining admissible validity. The purpose of this is to make the transition from *admissible validity* to *validity* in the argument for canonicity, as outlined in Section 3. Before formally stating the rules which can be applied at each stage of the algorithm, we outline its general strategy.

The final stage of the algorithm, where the actual elimination of propositional variables takes place, is the application of the two Ackermann rules. There are very specific syntactic requirements of the quasi-inequalities which can be used as input for the right and left Ackermann rules, (RA) and (LA). The steps before this final elimination stage are used to get the original inequality into the correct ‘shape’ so that the Ackermann rules can be applied.

- **Preprocessing:** operations are distributed according to the signed generation trees (see below), inequalities are split where possible with (\vee LA) and (\wedge RA), and simple propositional variable elimination takes place using (\perp) and (\top). Preprocessing may split the original inequality into a number of inequalities, on each of which we proceed separately.
- **Conversion of fixed point binders:** all occurrences of $\mu X.\varphi(X)$ are converted to $\mu^* X.\varphi(X)$ and all occurrences of $\nu X.\psi(X)$ are converted to $\nu^* X.\psi(X)$. We emphasize that this step is required in both tame *and* proper runs of μ^* -ALBA (see below).
- **First approximation:** an inequality is converted by (FA) into a quasi-inequality consisting of an implication with two inequalities in the antecedent, and one inequality in the consequent. The inequality in the consequent contains no propositional variables and thus all steps after this point are aimed only at eliminating propositional variables from the two inequalities in the antecedent.
- **Residuation, adjunction and approximation rules:** these rules prepare the antecedents of quasi-inequalities for the application of the Ackermann rules. The residuation and adjunction rules are straightforward applications of properties of the operations on perfect algebras. The approximation rules are more intricate in their formulation. The approximation rules for fixed point binders are the crucial link which enables us to apply this algorithmic approach to the mu-calculus. When dealing with inductive inequalities, a winning strategy (see Section 10) dictating the order in which these various rules are to be applied is determined by the way in which each critical branch is divided into its P_1, P_2 and P_3 parts as described in the previous section.
- **Ackermann rules:** the rules (RA) and (LA) are applied to the quasi-inequalities in the final step to eliminate the propositional variables.

Before we introduce the rules of our algorithm we comment briefly on their presentation. Let φ and ψ be two \mathcal{L}_1 -terms. Then $\varphi \leq \psi$ is an \mathcal{L}_1 -inequality. We use the symbols $\&$ and \Rightarrow , interpreted as conjunction and implication, respectively, to combine \mathcal{L} -inequalities into *quasi-inequalities*. To show canonicity of an inequality, we must show that the validity of the inequality is preserved when we move from interpreting it in an algebra \mathbf{A} to interpreting it in the canonical extension \mathbf{A}^δ . The algorithm μ^* -ALBA restructures the inequality into an equivalent quasi-inequality (or sometimes a *set of* quasi-inequalities).

As mentioned in the introduction, when formulas from the extended language \mathcal{L}^+ are interpreted in \mathbf{A}^δ , an assignment V will have $V(\text{NOM}) \subseteq J^\infty(\mathbf{A}^\delta)$ and $V(\text{CNOM}) \subseteq M^\infty(\mathbf{A}^\delta)$. Given an assignment V , we call V' a p -variant (or, potentially a **j**-variant or **m**-variant) of V if V' agrees with V on all elements of $\text{PROP} \cup \text{NOM} \cup \text{CNOM}$ except possibly at p (respectively, at **j** or **m**). We denote this by writing $V' \sim_p V$. From this point on we will use the same V to denote both the assignment $V : \text{PROP} \cup \text{NOM} \cup \text{CNOM} \rightarrow \mathbf{A}^\delta$ and the assignment $V : \mathcal{L}^+ \rightarrow \mathbf{A}^\delta$.

Preprocessing Given an inequality $\varphi \leq \psi$, consider the positive generation tree of φ and the negative generation tree of ψ and apply the following transformations exhaustively:

1. distribute positive occurrences of \diamond and \wedge over positive occurrences of \vee ;
2. distribute negative occurrences of \square and \vee over negative occurrences of \wedge ;
3. apply the *splitting rules* rules (\vee LA) and (\wedge RA);
4. apply the *monotone variable elimination rules* (\top) and (\perp).

In (\perp) we require that α is negative and β is positive in p , while for (\top) we require that γ is positive and χ is negative in p .

$$\frac{\alpha \vee \beta \leq \gamma}{\alpha \leq \gamma \quad \beta \leq \gamma} (\vee\text{LA}) \quad \frac{\alpha \leq \beta \wedge \gamma}{\alpha \leq \beta \quad \alpha \leq \gamma} (\wedge\text{RA}) \quad \frac{\alpha(p) \leq \beta(p)}{\alpha(\perp) \leq \beta(\perp)} (\perp) \quad \frac{\gamma(p) \leq \chi(p)}{\gamma(\top) \leq \chi(\top)} (\top)$$

First approximation rule This is the same as for μ -ALBA, viz.:

$$\frac{\varphi \leq \psi}{\forall \mathbf{j} \forall \mathbf{m}[(\mathbf{j} \leq \varphi \ \& \ \psi \leq \mathbf{m}) \Rightarrow \mathbf{j} \leq \mathbf{m}]} (\text{FA})$$

The rule (FA) takes an inequality from \mathcal{L}_* (or \mathcal{L}_1) and replaces it with a quasi-inequality from \mathcal{L}^+ (\mathcal{L}_1^+). The soundness of (FA) follows from the set $J^\infty(\mathbf{A}^\delta)$ (the interpretation of the nominals of \mathcal{L}^+) being join-dense in \mathbf{A}^δ , and from $M^\infty(\mathbf{A}^\delta)$ (the interpretation of the co-nominals of \mathcal{L}^+) being meet-dense in \mathbf{A}^δ (see the base case in the proof of Proposition 9.1).

Residuation rules These rules are the same as the corresponding group for μ -ALBA. The soundness and invertibility of the residuation rules follows from the facts that \vee and \wedge are, respectively, the right residual of $-$ and the left residual of \rightarrow .

$$\frac{\chi - \psi \leq \varphi}{\chi \leq \psi \vee \varphi} (-\text{LR}) \quad \frac{\varphi \leq \chi \rightarrow \psi}{\varphi \wedge \chi \leq \psi} (\rightarrow\text{RR}) \quad \frac{\chi \wedge \psi \leq \varphi}{\chi \leq \psi \rightarrow \varphi} (\wedge\text{LR}) \quad \frac{\varphi \leq \chi \vee \psi}{\varphi - \chi \leq \psi} (\vee\text{RR})$$

Adjunction rules The equivalent validity of the formulas above and below the line in (\vee LA) and (\wedge RA) follows from the fact that \vee is a left adjoint and \wedge a right adjoint of the diagonal map $a \mapsto (a, a)$. The rules (\diamond LA) and (\square RA) are justified by \diamond being the left adjoint of \blacksquare , and \square being the right adjoint of \blacklozenge .

$$\frac{\varphi \vee \chi \leq \psi}{\varphi \leq \psi \quad \chi \leq \psi} (\vee\text{LA}) \quad \frac{\psi \leq \varphi \wedge \chi}{\psi \leq \varphi \quad \psi \leq \chi} (\wedge\text{RA}) \quad \frac{\diamond \varphi \leq \psi}{\varphi \leq \blacksquare \psi} (\diamond\text{LA}) \quad \frac{\varphi \leq \square \psi}{\blacklozenge \varphi \leq \psi} (\square\text{RA})$$

Approximation rules These four rules have the requirement that the nominals and co-nominals introduced by them need to be *fresh*, i.e., do not occur in the derivation thus far. The justification of all four rules uses the join-primeness of elements of $J^\infty(\mathbf{A}^\delta)$ and the meet-primeness of elements of $M^\infty(\mathbf{A}^\delta)$ (see the example of $(\rightarrow \text{Appr})$ in Proposition 9.1).

$$\begin{array}{c} \frac{\Box\psi \leq \mathbf{m}}{\exists \mathbf{n}(\Box \mathbf{n} \leq \mathbf{m} \ \& \ \psi \leq \mathbf{n})} (\Box \text{Appr}) \quad \frac{\mathbf{j} \leq \Diamond\psi}{\exists \mathbf{i}(\mathbf{j} \leq \Diamond \mathbf{i} \ \& \ \mathbf{i} \leq \psi)} (\Diamond \text{Appr}) \\[10pt] \frac{\chi \rightarrow \varphi \leq \mathbf{m}}{\exists \mathbf{j} \exists \mathbf{n}(\mathbf{j} \rightarrow \mathbf{n} \leq \mathbf{m} \ \& \ \mathbf{j} \leq \chi \ \& \ \varphi \leq \mathbf{n})} (\rightarrow \text{Appr}) \quad \frac{\mathbf{i} \leq \chi - \varphi}{\exists \mathbf{j} \exists \mathbf{n}(\mathbf{i} \leq \mathbf{j} - \mathbf{n} \ \& \ \mathbf{j} \leq \chi \ \& \ \varphi \leq \mathbf{n})} (-\text{Appr}) \end{array}$$

We will usually refer to the above group of approximation rules as the *ordinary* approximation rules in order to distinguish them from the fixed point binder approximation rules that follow.

Approximation rules for fixed point binders Our next two rules are used to extract propositional variables that occur within the scope of a fixed point binder. Let ψ be an $(n+1)$ -ary term function, and suppose that there exists an order type τ over n such that $\psi(\bar{x}, X)$ is completely \vee -preserving in $(\bar{x}, X) \in \mathbf{C}^\tau \times \mathbf{C}$ where \mathbf{C} is any perfect mu-algebra of the second kind. Then (subject to the conditions listed below) we will be able to define two rules: $(\mu^\tau\text{-A-R})$ and $(\nu^\tau\text{-A-R})$. We refer the reader to the definitions at the end of Section 2 regarding order types and join-irreducible elements to assist in the reading of these two rules.

$$\begin{array}{c} \frac{\mathbf{i} \leq \mu^* X. \psi(\bar{\varphi}/\bar{x}, X)}{\bigvee_{i=1}^n (\exists \mathbf{j}^{\tau_i} [\mathbf{i} \leq \mu^* X. \psi(\bar{\mathbf{j}}_i^\tau / \bar{x}, X) \ \& \ \mathbf{j}^{\tau_i} \leq \varphi_i])} (\mu^\tau\text{-A-R}) \\[10pt] \frac{\nu^* X. \varphi(\bar{\psi}/\bar{x}, X) \leq \mathbf{m}}{\bigvee_{i=1}^n (\exists \mathbf{n}^{\tau_i} [\nu^* X. \varphi(\bar{\mathbf{n}}_i^\tau / \bar{x}, X) \leq \mathbf{m} \ \& \ \psi_i \leq \mathbf{n}^{\tau_i}])} (\nu^\tau\text{-A-R}) \end{array}$$

where

1. in each rule, the variables $\bar{x} \in \text{PHVAR}$ do not occur in any formula in $\bar{\psi}$ or in $\bar{\varphi}$;
2. all propositional variables and free fixed point variables in $\psi(\bar{x}, X)$ and $\varphi(\bar{x}, X)$ are among \bar{x} and X .
3. in $(\mu^\tau\text{-A-R})$ the associated term function of $\psi(\bar{x}, X)$ is completely \vee -preserving in $(\bar{x}, X) \in \mathbf{C}^\tau \times \mathbf{C}$, for any perfect modal bi-Heyting algebra \mathbf{C} ; in particular we require that $\psi(\bar{x}, X)$ is positive (negative) in x_i if $\tau_i = 1$ ($\tau_i = \partial$);
4. in $(\nu^\tau\text{-A-R})$ the associated term function of $\varphi(\bar{x}, X)$ is completely \wedge -preserving in $(\bar{x}, X) \in \mathbf{C}^\tau \times \mathbf{C}$, for any perfect modal bi-Heyting algebra \mathbf{C} ; in particular we require that $\varphi(\bar{x}, X)$ is positive (negative) in x_i if $\tau_i = 1$ ($\tau_i = \partial$).

Note that the difference between these rules and the more general rules, $(\mu^\tau\text{-A})$ and $(\nu^\tau\text{-A})$ (see [5, Section 2.3]), is the absence of the additional tuple, \bar{z} , of place holder variables in φ and ψ . This ensures that, in applications of the rule, the resulting inequalities $\mathbf{i} \leq \mu^* X. \psi(\bar{\mathbf{j}}_i^\tau / \bar{x}, X)$ and $\nu^* X. \varphi(\bar{\mathbf{n}}_i^\tau / \bar{x}, X) \leq \mathbf{m}$ are pure — a fact which will be needed in Section 10 to justify the applicability of the Ackermann-rules, when we prove that $\mu^*\text{-ALBA}$ successfully purifies all restricted and tame inductive mu-inequalities.

Note that $(\mu^\tau\text{-A-R})$ and $(\nu^\tau\text{-A-R})$ are the only rules which can result in the overall quasi-inequality becoming a set of quasi-inequalities. This happens as a result of the introduction of the disjunction $\bigvee_{i=1}^n$. Given the way we have set up $\mu^*\text{-ALBA}$, when the rule is applied, the inequality $\mathbf{i} \leq \mu^* X. \psi(\bar{\varphi}/\bar{x}, X)$ will occur in the antecedent of an implication as part of a conjunction of inequalities. Once the disjunction $\bigvee_{i=1}^n$ is introduced, this will be distributed over the conjunctions and the implications, producing a set of n quasi-inequalities.

Example 5.1. Consider the mu-inequality $\mu^*X.(p \vee (\Box \perp - q) \vee \Diamond X) \leq \Diamond \Box (p \vee q)$. The first approximation rule (FA) gives us the quasi-inequality $\forall \mathbf{i} \forall \mathbf{m} [\mathbf{i} \leq \mu^*X.(p \vee (\Box \perp - q) \vee \Diamond X) \ \& \ \Diamond \Box (p \vee q) \leq \mathbf{m} \Rightarrow \mathbf{i} \leq \mathbf{m}]$. The first inequality from the antecedent gives us the term function $\psi(x, y, X) = x \vee (\Box \perp - y) \vee \Diamond X$ which is \vee -preserving in (x, y, X) for the order type $\tau = (1, \partial)$. That is, $\psi(x, y, X)$ is order-preserving as a map from $\mathbf{C}^\tau \times \mathbf{C}$. A tuple $\bar{\mathbf{j}}_1^\tau$ has the form (\mathbf{j}, \top) and a tuple $\bar{\mathbf{j}}_2^\tau$ has the form (\perp, \mathbf{n}) . Let $\bar{\varphi} = (p, q)$. The rule $(\mu^\tau\text{-A-R})$ gives us the following:

$$\frac{\mathbf{i} \leq \mu^*X.\psi(\bar{\varphi}/\bar{x}, X)}{(\exists \mathbf{j}^1 [\mathbf{i} \leq \mu^*X.\psi(\bar{\mathbf{j}}_1^\tau/\bar{x}, X) \ \& \ (\mathbf{j}^1 \leq^1 p)] \ \wp \ (\exists \mathbf{j}^\partial [\mathbf{i} \leq \mu^*X.\psi(\bar{\mathbf{j}}_2^\tau/\bar{x}, X) \ \& \ (\mathbf{j}^\partial \leq^\partial q)])}$$

This can be further simplified and written as:

$$\frac{\mathbf{i} \leq \mu^*X.\psi(\bar{\varphi}/\bar{x}, X)}{(\exists \mathbf{j} [\mathbf{i} \leq \mu^*X.(\mathbf{j} \vee (\Box \perp - \top) \vee X) \ \& \ (\mathbf{j} \leq p)] \ \wp \ (\exists \mathbf{n} [\mathbf{i} \leq \mu^*X.(\perp \vee (\Box \perp - \mathbf{n}) \vee X) \ \& \ (\mathbf{n} \geq q)])}$$

Ackermann rules The Ackermann rules are used for the crucial task of eliminating propositional variables from quasi-inequalities. Before we define these rules we need to introduce some terminology and definitions.

Let $\varphi \in \mathcal{L}_1^+$ or $\varphi \in \mathcal{L}_*^+$. We say that φ is *positive (negative)* in a variable p if in the generation tree $+\varphi$ all p -nodes are signed $+$ ($-$). An inequality $\varphi \leq \psi$ is *positive (negative)* in a variable p if φ is negative (positive) in p and ψ is positive (negative) in p .

Definition 5.2. The *syntactically open formulas* φ and *syntactically closed formulas* ψ are defined by simultaneous mutual recursion as follows:

$$\begin{aligned} \varphi &::= \perp \mid \top \mid p \mid \mathbf{m} \mid \varphi_1 \wedge \varphi_2 \mid \varphi_1 \vee \varphi_2 \mid \psi \rightarrow \varphi \mid \varphi - \psi \mid \Box \varphi \mid \Diamond \varphi \mid \blacksquare \varphi \mid \nu^* X.\varphi \\ \psi &::= \perp \mid \top \mid p \mid \mathbf{i} \mid \psi_1 \wedge \psi_2 \mid \psi_1 \vee \psi_2 \mid \varphi \rightarrow \psi \mid \psi - \varphi \mid \Box \psi \mid \Diamond \psi \mid \blacklozenge \psi \mid \mu^* X.\psi \end{aligned}$$

where $p \in \text{PROP}$, $\mathbf{i} \in \text{NOM}$, and $\mathbf{m} \in \text{CNOM}$.

The *syntactically almost open formulas* and *syntactically almost closed formulas* are defined by adding, respectively, $\mu^*X.\varphi$ and $\nu^*X.\psi$ to the recursions above.

Informally, an \mathcal{L}_*^+ -term is *syntactically almost open* if, in it, all occurrences of nominals and \blacklozenge are negative, while all occurrences of co-nominals and \blacksquare are positive. If, in addition, all occurrences of μ^* are negative and all occurrences of ν^* positive, the term is syntactically open. Similarly, an \mathcal{L}_*^+ -term is *syntactically almost closed* if, in it, all occurrences of nominals and \blacklozenge are positive, while all occurrences of co-nominals and \blacksquare are negative. If, in addition, all occurrences of μ^* are positive and all occurrences of ν^* are negative, the term is syntactically closed.

Given these definitions we are now able to define our restricted versions of the Ackermann rules:

$$\frac{\exists p [\&_{i=1}^n \alpha_i \leq p \ \& \ \&_{j=1}^m \beta_j(p) \leq \gamma_j(p)]}{\&_{j=1}^m \beta_j(\bigvee_{i=1}^n \alpha_i / p) \leq \gamma_j(\bigvee_{i=1}^n \alpha_i / p)} \text{ (RA)}$$

subject to the restrictions that the α_i are p -free and syntactically closed, the β_j are positive in p and syntactically closed, while the γ_j are negative in p and syntactically open.

$$\frac{\exists p [\&_{i=1}^n p \leq \alpha_i \ \& \ \&_{j=1}^m \gamma_j(p) \leq \beta_j(p)]}{\&_{j=1}^m \gamma_j(\bigwedge_{i=1}^n \alpha_i / p) \leq \beta_j(\bigwedge_{i=1}^n \alpha_i / p)} \text{ (LA)}$$

subject to the restrictions that the α_i are p -free and syntactically open, the β_j are positive in p and syntactically open, while the γ_j are negative in p and syntactically closed.

A *tame run* of μ^* -ALBA is one during which there are no applications of either $(\mu^\tau\text{-A-R})$ or $(\nu^\tau\text{-A-R})$. By contrast, a *proper run* of μ^* -ALBA is one during which *all* occurrences of fixed point binders are handled by $(\mu^\tau\text{-A-R})$ and $(\nu^\tau\text{-A-R})$. We say that a run of the algorithm μ^* -ALBA *succeeds* if all propositional variables are eliminated from the input inequality, $\varphi \leq \psi$, and we denote the resulting set of pure quasi-inequalities by $\text{pure}(\varphi^* \leq \psi^*)$. An inequality on which some run of μ^* -ALBA succeeds is called a μ^* -ALBA *inequality*.

6 Syntactic conditions for meet and join preservation: inner formulas

In the formulation of the approximation rules $(\mu^\tau\text{-A-R})$ and $(\nu^\tau\text{-A-R})$ we required the term functions $\psi(\bar{x}, X)$ and $\varphi(\bar{x}, X)$ to be, respectively, completely \vee and \wedge -preserving as maps from $\mathbf{A}^\tau \times \mathbf{A}$ to \mathbf{A} . The inner formulas (introduced in [5, Section 4]) are a syntactically specified classes of formulas, the term functions of which satisfy these properties. Being an inner formula is thus an effectively checkable sufficient condition for the applicability of the rules $(\mu^\tau\text{-A-R})$ and $(\nu^\tau\text{-A-R})$.

Definition 6.1. Let $\bar{y}, \bar{z} \subseteq \text{PHVAR}$ and $\bar{X} \subseteq \text{FVAR}$ be tuples of variables which are pairwise different in the union of the their underlying sets. Let τ be an order-type on $\bar{x} = \bar{y} \oplus \bar{X}$. The τ - \square and τ - \diamond (\bar{x}, \bar{z}) -inner formulas $((\bar{x}, \bar{z})\text{-IF}_\tau^\square$ and $(\bar{x}, \bar{z})\text{-IF}_\tau^\diamond$), the free variables of which are contained in (\bar{x}, \bar{z}) , are given by the following simultaneous recursion (for the sake of readability, the parameters \bar{x} and \bar{z} are omitted):

$$\begin{aligned} \text{IF}_\tau^\square \ni \varphi &::= x_i \mid \square\varphi \mid \varphi_1 \wedge \varphi_2 \mid \nu^* Y.\varphi' \mid \pi \rightarrow \varphi \mid \pi \vee \varphi \mid \psi^c \rightarrow \pi \\ \text{IF}_\tau^\diamond \ni \psi &::= x_i \mid \diamond\psi \mid \psi_1 \vee \psi_2 \mid \mu^* Y.\psi' \mid \psi - \pi \mid \pi \wedge \psi \mid \pi - \varphi^c \end{aligned}$$

where

1. $\tau_i = 1$ in the base of the recursion,
2. π is $\pi(\bar{z}) \in \mathcal{L}_*$ (specifically, $\pi(\bar{z})$ contains none of the variables in \bar{x} or \bar{X}),
3. $\varphi' = \varphi'(\bar{y} \oplus \bar{X}', \bar{z})$ and $\psi' = \psi'(\bar{y} \oplus \bar{X}', \bar{z})$ are $\text{IF}_{\tau'}^\square$ and $\text{IF}_{\tau'}^\diamond$, respectively, with $\bar{X}' = \bar{X} \oplus Y$ and $\tau' = \tau \oplus 1$,
4. $\psi^c \in (\bar{x}, \bar{z})\text{-IF}_{\tau'}^\diamond$ and $\varphi^c \in (\bar{x}, \bar{z})\text{-IF}_{\tau'}^\square$.
5. All other formulas have their free variables among (\bar{x}, \bar{z}) .

The key fact about $(\bar{x}, \bar{z})\text{-IF}_\tau^\square$ and $(\bar{x}, \bar{z})\text{-IF}_\tau^\diamond$ formulas is the following:

Lemma 6.2 ([5, Lemma 4.3]). *For any perfect modal bi-Heyting algebra \mathbf{A} , the term function associated with any IF_τ^\square formula $\varphi(\bar{x}, \bar{z})$ (resp., IF_τ^\diamond formula $\psi(\bar{x}, \bar{z})$) is completely meet-preserving (resp., join-preserving) as a map $\mathbf{A}^\tau \rightarrow \mathbf{A}$, fixing the variables \bar{z} .*

In particular, if the \bar{y} are constant \mathcal{L}_^+ sentences, then the term function associated with $\varphi(\bar{x}, \bar{y}/\bar{z})$ (resp., $\psi(\bar{x}, \bar{y}/\bar{z})$) is completely meet-preserving (resp., join-preserving) as a map $\mathbf{A}^\tau \rightarrow \mathbf{A}$.*

For our purposes then, the most important consequence of this lemma is the fact that $(\nu^\tau\text{-A-R})$ and $(\mu^\tau\text{-A-R})$ are respectively applicable to formulas of the form $\psi(\bar{\varphi}/\bar{x}, \bar{y}/\bar{z})$ and $\varphi(\bar{\psi}/\bar{x}, \bar{y}/\bar{z})$, where $\psi(\bar{x}, \bar{z})$ is IF_τ^\square , $\psi(\bar{x}, \bar{z})$ is IF_τ^\diamond and the \bar{y} are constant \mathcal{L}_*^+ sentences. In particular, $(\nu^\tau\text{-A-R})$ and $(\mu^\tau\text{-A-R})$ are applicable to formulas $\psi(\bar{\varphi}/\bar{x})$ and $\varphi(\bar{\psi}/\bar{x})$, where $\psi(\bar{x}, \bar{z})$ is IF_τ^\square and $\psi(\bar{x}, \bar{z})$ is IF_τ^\diamond with \bar{z} the empty tuple.

7 Soundness of the fixed point approximation rules

Before we prove the soundness of the Approximation Rules (μ^τ -A-R) and (ν^τ -A-R), we require some lemmas regarding preservation properties of operations.

Lemma 7.1. *Let \mathbf{L} be a complete lattice. If $f : \mathbf{L}^m \times \mathbf{L} \rightarrow \mathbf{L}$ is completely join-preserving, then the function $f(x_1, \dots, x_m, \perp) : \mathbf{L}^m \rightarrow \mathbf{L}$ is completely join-preserving.*

Proof. Let $S \subseteq \mathbf{L}^m$. We want to show that $f(\bigvee S, \perp) = \bigvee \{f(s, \perp) \mid s \in S\}$. Let $S' = \{(s, \perp) \mid s \in S\} \subseteq \mathbf{L} \times \mathbf{L}^m$. Then $\bigvee S' = (\bigvee S, \perp)$ and $f(\bigvee S') = \bigvee \{f(s') \mid s' \in S'\} = \bigvee \{f(s, \perp) \mid s \in S\}$. \square

We note that this result would also hold if we replaced \perp with any $a \in L$, so long as S is non-empty. To accommodate the case that $S = \emptyset$, we must have $a = \perp$.

Lemma 7.2. *Let \mathbf{A} be a mu-algebra of the second kind and let τ be an order type over n . Let τ' be the order type over $n+1$ defined by $\tau' = \tau \oplus \{1\}$. Suppose $f : \mathbf{A}^{n+1} \rightarrow \mathbf{A}$ is an \mathcal{L}_* term function such that $f^{\mathbf{A}^\delta} : (\mathbf{A}^\delta)^{\tau'} \rightarrow \mathbf{A}^\delta$ is completely \bigvee -preserving. By the assumption that \mathbf{A} is of the second kind, $\mu_2 x. f(a_1, a_2, \dots, a_n, x)$ exists in \mathbf{A} for all $a_1, a_2, \dots, a_n \in A$.*

Let $S \subseteq (\mathbf{A}^\delta)^\tau$ such that $\bigvee S \in \mathbf{A}^\tau$. Then $\mu^ x. f^{\mathbf{A}^\delta}(\bigvee S, x) = \bigvee \{\mu^* x. f^{\mathbf{A}^\delta}(s, x) \mid s \in S\}$.*

Proof. We have the following sequence of equalities:

$$\mu^* x. f^{\mathbf{A}^\delta}(\bigvee S, x) = \bigwedge \{a \in A \mid f^{\mathbf{A}^\delta}(\bigvee S, a) \leq a\} \quad (1)$$

$$= \bigwedge \{a \in A \mid f^{\mathbf{A}}(\bigvee S, a) \leq a\} \quad (2)$$

$$= \mu x. f^{\mathbf{A}}(\bigvee S, x) \quad (3)$$

$$= \mu_2 x. f^{\mathbf{A}}(\bigvee S, x) \quad (4)$$

$$= \bigvee_{\alpha \geq 0} (f^{\mathbf{A}})^\alpha(\bigvee S, \perp) \quad (5)$$

$$= \bigvee_{\alpha \geq 0} (f^{\mathbf{A}^\delta})^\alpha(\bigvee S, \perp). \quad (6)$$

The equivalence of (1) and (2), as well as (5) and (6), follows from the fact that all of the arguments of $f^{\mathbf{A}^\delta}$ are in \mathbf{A} . The equivalence of (2) and (3) is the definition of how $\mu x. \varphi(x)$ is interpreted, and the equivalence of (3) and (4) follows from the fact that \mathbf{A} is a mu-algebra of the second kind.

By induction on α we will show that $(f^{\mathbf{A}^\delta})^\alpha(\bigvee S, \perp) = \bigvee \{(f^{\mathbf{A}^\delta})^\alpha(s, \perp) \mid s \in S\}$ and this will be sufficient to prove the overall result.

- Case $\alpha = 0$: $(f^{\mathbf{A}^\delta})^0(\bigvee S, \perp) = \perp$.
- Case $\alpha = 1$: $(f^{\mathbf{A}^\delta})^1(\bigvee S, \perp) = f^{\mathbf{A}^\delta}(\bigvee S, \perp) = \bigvee \{f^{\mathbf{A}^\delta}(s, \perp) \mid s \in S\}$. The second equivalence follows from Lemma 7.1.
- Successor ordinals: $(f^{\mathbf{A}^\delta})^{\alpha+1}(\bigvee S, \perp) = f^{\mathbf{A}^\delta}(\bigvee S, (f^{\mathbf{A}^\delta})^\alpha(\bigvee S, \perp)) = f^{\mathbf{A}^\delta}(\bigvee S, \bigvee \{(f^{\mathbf{A}^\delta})^\alpha(s, \perp) \mid s \in S\})$ by the inductive hypothesis. Since $f^{\mathbf{A}^\delta}$ is completely join-preserving we have

$$f^{\mathbf{A}^\delta}(\bigvee S, \bigvee_{s \in S} (f^{\mathbf{A}^\delta})^\alpha(s, \perp)) = \bigvee_{s \in S} f^{\mathbf{A}^\delta}(s, (f^{\mathbf{A}^\delta})^\alpha(s, \perp)) = \bigvee_{s \in S} (f^{\mathbf{A}^\delta})^{\alpha+1}(s, \perp).$$

- Limit ordinals:

$$\begin{aligned} (f^{\mathbf{A}^\delta})^\gamma(\bigvee S, \perp) &= \bigvee_{\beta < \gamma} (f^{\mathbf{A}^\delta})^\beta(\bigvee S, \perp) = \bigvee_{\beta < \gamma} \bigvee_{s \in S} (f^{\mathbf{A}^\delta})^\beta(s, \perp) \quad (\text{by IH}) \\ &= \bigvee_{s \in S} \bigvee_{\beta < \gamma} (f^{\mathbf{A}^\delta})^\beta(s, \perp) = \bigvee_{s \in S} (f^{\mathbf{A}^\delta})^\gamma(s, \perp). \end{aligned} \quad \square$$

Before demonstrating that the Approximation Rule is sound we should point out that this rule is only ever applied to a *quasi-inequality* that is the result of an application of the First Approximation Rule (FA). That is, $(\mu^\tau\text{-A-R})$ and $(\nu^\tau\text{-A-R})$ are each applied to an inequality which forms *part of* the antecedent of an implication. When demonstrating the soundness of the rule $(\mu^\tau\text{-A-R})$ it is therefore sufficient to show that the inequality above the line and the inequality below the line are valid under assignments which agree everywhere except at some nominal which does not occur in the consequent of the quasi-inequality.

Proposition 7.3. (Soundness of $(\mu^\tau\text{-A-R})$) *Let \mathbf{C} be a mu-algebra of the second kind. Let $\psi(\bar{x}, X)$ and $\bar{\varphi}$ be terms in \mathcal{L}_* , with $\psi(\bar{x}, X)$ completely \vee -preserving in $(\bar{x}, X) \in (\mathbf{C}^\delta)^\tau \times \mathbf{C}^\delta$ for an order type τ . Let V be an admissible assignment on \mathbf{C}^δ .*

Then $\mathbf{C}^\delta, V \models \mathbf{i} \leq \mu^ X. \psi(\bar{\varphi}/\bar{x}, X)$ if and only if there exists i ($1 \leq i \leq n$) and a \mathbf{j}^{τ_i} -variant V' of V such that*

$$\mathbf{C}^\delta, V' \models \mathbf{i} \leq \mu^* X. \psi(\bar{\mathbf{j}}_i^\tau / \bar{x}, X) \quad \text{and} \quad \mathbf{C}^\delta, V' \models \mathbf{j}^{\tau_i} \leq^{\tau_i} \varphi_i.$$

Proof. Suppose that $\mathbf{C}^\delta, V \models \mathbf{i} \leq \mu^* X. \psi(\bar{\varphi}/\bar{x}, X)$ for some admissible assignment V . By Lemma 7.2 we have that the term function $\mu^* X. \psi(\bar{x}, X)$ is completely \vee -preserving in $(\mathbf{C}^\delta)^\tau$.

At this point we will not distinguish between formulas and their interpretations under V . Since \mathbf{C}^δ is a perfect modal bi-Heyting algebra we have that $\bar{\varphi} = \bigvee \{ \bar{j} \in J^\infty((\mathbf{C}^\delta)^\tau) \mid \bar{j} \leq \bar{\varphi} \}$. Thus we have

$$\mu^* X. \psi(\bar{\varphi}, X) = \mu^* X. \psi(\bigvee \bar{j}, X) = \bigvee \{ \mu^* X. \psi(\bar{j}, X) \mid \bar{j} \leq \bar{\varphi} \}.$$

Now $V(\mathbf{i}) \leq \bigvee \{ \mu^* X. \psi(\bar{j}, X) \mid \bar{j} \leq \bar{\varphi} \}$. Since $V(\mathbf{i})$ is completely join-irreducible and hence completely join-prime we have that there exists $\bar{j}_0 \in J^\infty((\mathbf{C}^\delta)^\tau)$ with $\bar{j}_0 \leq \bar{\varphi}$ such that $V(\mathbf{i}) \leq \mu^* X. \psi(\bar{j}_0, X)$. Recall from the end of Section 2 that \bar{j}_0 is \perp^{τ_i} at every $1 \leq i \leq n$ except at one coordinate, say k . There we have $(\bar{j}_0)_k \in J^\infty((\mathbf{C}^\delta)^{\tau_k})$ and also $(\bar{j}_0)_k \leq^{\tau_k} \bar{\varphi}_k$. Let \mathbf{j}_0 be some nominal for which $\mu^* X. \psi(\bar{\varphi}/\bar{x}, X)$ is \mathbf{j}_0 -free. Now let V' be the \mathbf{j}_0 -variant of V such that $V'(\mathbf{j}_0) = (\bar{j}_0)_k$.

For the converse, suppose that there exists $i \in \{1, \dots, n\}$ and \mathbf{j}^{τ_i} such that $\mathbf{C}^\delta, V' \models \mathbf{i} \leq \mu^* X. \psi(\bar{\mathbf{j}}_i^\tau / \bar{x}, X)$ and $\mathbf{C}^\delta, V' \models \mathbf{j}^{\tau_i} \leq^{\tau_i} \varphi_i$ where V' is an admissible \mathbf{j}^{τ_i} -variant of V . Let us consider the n -tuple $\bar{\mathbf{j}}_i^\tau$. Given $k \neq i$, if $\tau_k = 1$, then the k -th coordinate of $\bar{\mathbf{j}}_i^\tau$ is \perp and hence $(\bar{\mathbf{j}}_i^\tau)_k \leq \varphi_k$. Again for $k \neq i$, if $\tau_k = \partial$, then the k -th coordinate of $\bar{\mathbf{j}}_i^\tau$ is \top and thus $(\bar{\mathbf{j}}_i^\tau)_k \leq^\partial \varphi_k$. Thus we have for all k that $(\bar{\mathbf{j}}_i^\tau)_k \leq^{\tau_k} \varphi_k$. The fact that $\mathbf{i} \leq \mu^* X. \psi(\bar{\varphi}/\bar{x}, X)$ follows from the fact that ψ is completely join-preserving (and hence monotone) in $\bar{x} \in (\mathbf{C}^\delta)^\tau$. Thus we have

$$V'(\mathbf{i}) \leq V'(\mu^* X. \psi(\bar{\mathbf{j}}_i^\tau / \bar{x}, X)) \leq V'(\mu^* X. \psi(\bar{\varphi}/\bar{x}, X))$$

and hence $\mathbf{C}^\delta, V \models \mathbf{i} \leq \mu^* X. \psi(\bar{\varphi}/\bar{x}, X)$. □

The statements and proofs of Lemma 7.1, Lemma 7.2 and Proposition 7.3 can easily be dualised and hence we can prove the soundness of the rule $(\nu^\tau\text{-A-R})$.

The reason that we need to use μ^* and ν^* in the rules $(\mu^\tau\text{-A-R})$ and $(\nu^\tau\text{-A-R})$, respectively, is shown by the proof of Lemma 7.2. If we were to use $\mu x. f^{\mathbf{A}^\delta}(\bigvee S, x) = \bigwedge \{ a \in A^\delta \mid f^{\mathbf{A}^\delta}(\bigvee S, a) \leq a \}$ in line (1), we would then have only $(1) \leq (2)$ (as $A \subseteq A^\delta$). Thus if we formulated $(\mu^\tau\text{-A-R})$ and Lemma 7.2 with μ instead of μ^* , we would not have the equality in Lemma 7.2 and thus would not be able to show the invariance of admissible validity under $(\mu^\tau\text{-A-R})$.

8 Soundness of the Ackermann rules

In this section we prove the soundness of the Ackermann rules, (RA) and (LA). Once we have shown the soundness of (RA) and (LA), we will be able, in the next section, to prove Proposition 9.1. This proposition declares the soundness of μ^* -ALBA derivations with respect to admissible validity.

We will need the next two technical lemmas, the proofs of which will make extensive use of the algebraic results presented in the appendix. Our strategy closely follows that in [6]. However, we work algebraically whereas [6] proceeds in the setting of general frames. Moreover, we need to accommodate fixed point binders which are absent in [6].

Fix a modal bi-Heyting algebra \mathbf{A} of the first kind. The set of open elements of \mathbf{A}^δ , denoted $\mathbb{O}(\mathbf{A}^\delta)$, is defined as $\{\bigvee S \mid S \subseteq A\}$, i.e., as all those elements of \mathbf{A}^δ that can be obtained as arbitrary joins of elements of \mathbf{A} . Dually, the set of closed elements of \mathbf{A}^δ , denoted $\mathbb{K}(\mathbf{A}^\delta)$, is defined as $\{\bigwedge S \mid S \subseteq A\}$, i.e., as all those elements of \mathbf{A}^δ that can be obtained as arbitrary meets of elements of \mathbf{A} . The intention behind the definition of syntactically open and closed \mathcal{L}_*^+ formulas is that admissible assignments will always interpret them as open and closed elements of \mathbf{A}^δ , respectively.

Lemma 8.1. *Let $\varphi(p, \bar{q}, \bar{\mathbf{i}}, \bar{\mathbf{m}})$ be syntactically closed and $\psi(p, \bar{q}, \bar{\mathbf{i}}, \bar{\mathbf{m}})$ syntactically open. Let $\bar{b} \in A$, $\bar{c} \in J^\infty(\mathbf{A}^\delta)$ and $\bar{d} \in M^\infty(\mathbf{A}^\delta)$. Let $k \in \mathbb{K}(\mathbf{A}^\delta)$ and $u \in \mathbb{O}(\mathbf{A}^\delta)$. Then,*

1. (a) *If $\varphi(p, \bar{q}, \bar{\mathbf{i}}, \bar{\mathbf{m}})$ is positive in p , then $\varphi(k, \bar{b}, \bar{c}, \bar{d}) \in \mathbb{K}(\mathbf{A}^\delta)$ and*
 (b) *if $\psi(p, \bar{q}, \bar{\mathbf{i}}, \bar{\mathbf{m}})$ is negative in p , then $\psi(k, \bar{b}, \bar{c}, \bar{d}) \in \mathbb{O}(\mathbf{A}^\delta)$.*
2. (a) *If $\varphi(p, \bar{q}, \bar{\mathbf{i}}, \bar{\mathbf{m}})$ is negative in p , then $\varphi(u, \bar{b}, \bar{c}, \bar{d}) \in \mathbb{K}(\mathbf{A}^\delta)$, and*
 (b) *if $\psi(p, \bar{q}, \bar{\mathbf{i}}, \bar{\mathbf{m}})$ is positive in p , then $\psi(u, \bar{b}, \bar{c}, \bar{d}) \in \mathbb{O}(\mathbf{A}^\delta)$.*

Proof. We proceed by simultaneous structural induction on φ and ψ . We show (1). Assume that $\varphi(p, \bar{q}, \bar{\mathbf{i}}, \bar{\mathbf{m}})$ is positive in p and $\psi(p, \bar{q}, \bar{\mathbf{i}}, \bar{\mathbf{m}})$ is negative in p . As they do not impact the overall result, we will omit the parameters \bar{q} , $\bar{\mathbf{i}}$ and $\bar{\mathbf{m}}$ and simply write $\varphi(p)$ and $\psi(p)$ for $\varphi(p, \bar{q}, \bar{\mathbf{i}}, \bar{\mathbf{m}})$ and $\psi(p, \bar{q}, \bar{\mathbf{i}}, \bar{\mathbf{m}})$ respectively. The base cases are when φ is of the form \top, \perp, p, q (where q is a propositional variable different from p), or $\varphi = \mathbf{i}$, and when ψ is of the form \top, \perp, q (where q is a propositional variable different from p), or $\psi = \mathbf{m}$. The φ cannot be a co-nominal \mathbf{m} since φ is syntactically closed but \mathbf{m} syntactically open. Similarly ψ cannot be p or a nominal \mathbf{i} since ψ must be negative in p and any occurrence of a nominal must be negative.

Now clearly \top, \perp, q are all interpreted as clopen elements of \mathbf{A}^δ , as is p . Furthermore, the claims follow for φ any nominal \mathbf{i} and ψ any co-nominal \mathbf{m} since $J^\infty(\mathbf{A}^\delta) \subseteq \mathbb{K}(\mathbf{A}^\delta)$ and $M^\infty(\mathbf{A}^\delta) \subseteq \mathbb{O}(\mathbf{A}^\delta)$.

If $\varphi(p) = \varphi_1(p) \wedge \varphi_2(p)$ or $\varphi(p) = \varphi_1(p) \vee \varphi_2(p)$ then both $\varphi_1(p)$ and $\varphi_2(p)$ must be syntactically closed and positive in p . Thus by the inductive hypothesis we have that $\varphi_1(k) \in \mathbb{K}(\mathbf{A}^\delta)$ and $\varphi_2(k) \in \mathbb{K}(\mathbf{A}^\delta)$ and both their meet and join are in $\mathbb{K}(\mathbf{A}^\delta)$ as this is closed under the lattice operations.

If $\psi(p) = \psi_1(p) \wedge \psi_2(p)$ or $\psi(p) = \psi_1(p) \vee \psi_2(p)$ then both $\psi_1(p)$ and $\psi_2(p)$ must be syntactically open and negative in p . By the inductive hypothesis, we have $\psi_1(k) \in \mathbb{O}(\mathbf{A}^\delta)$ and $\psi_2(k) \in \mathbb{O}(\mathbf{A}^\delta)$ and both their meet and join will be in $\mathbb{O}(\mathbf{A}^\delta)$ as this is also closed under the lattice operations.

We note that there cannot be occurrences of μ or ν in either $\varphi(p)$ or $\psi(p)$. If $\varphi(p)$ is of the form $\mu^*X.\varphi_1(p, X)$, then $\varphi(k) = \mu^*X.\varphi_1(k, X) = \bigwedge \{a \in A \mid \varphi_1(k, a) \leq a\}$ and so $\varphi(k) \in \mathbb{K}(\mathbf{A}^\delta)$. If $\psi(p)$ is of the form $\nu^*X.\psi_1(p, X)$, then $\psi(k) = \nu^*X.\psi_1(k, X) = \bigvee \{a \in A \mid a \leq \psi_1(k, a)\}$ and so $\psi(k) \in \mathbb{O}(\mathbf{A}^\delta)$.

If $\varphi(p) = \varphi_1(p) - \varphi_2(p)$ then $\varphi_1(p)$ is syntactically closed and positive in p and $\varphi_2(p)$ is syntactically open and negative in p . By the inductive hypothesis, $\varphi_1(k) \in \mathbb{K}(\mathbf{A}^\delta)$ and $\varphi_2(k) \in \mathbb{O}(\mathbf{A}^\delta)$. Now by Lemma .11(4) we have that $\varphi_1(k) - \varphi_2(k) = \varphi(k) \in \mathbb{K}(\mathbf{A}^\delta)$.

If $\psi(p) = \psi_1(p) \rightarrow \psi_2(p)$ then $\psi_1(p)$ is syntactically closed and positive in p while $\varphi_2(p)$ is syntactically open and negative in p . By the inductive hypothesis we have $\psi_1(k) \in \mathbb{K}(\mathbf{A}^\delta)$ and $\psi_2(k) \in \mathbb{O}(\mathbf{A}^\delta)$. Using Lemma .11(3) we see that $\psi_1(k) \rightarrow \psi_2(k) = \psi(k) \in \mathbb{O}(\mathbf{A}^\delta)$.

Now we look at the cases for the unary connectives. We note that φ cannot be of the form $\blacksquare\varphi_1$ as this is not syntactically closed. Likewise, ψ cannot be of the form $\blacklozenge\psi_1$ as this is not syntactically open.

If $\varphi(p)$ is of the form $\Box\varphi_1(p)$, $\Diamond\varphi_1(p)$ or $\blacklozenge\varphi_1(p)$ then $\varphi_1(p)$ must be syntactically closed and positive in p . By the inductive hypothesis we have that $\varphi_1(k) \in \mathbb{K}(\mathbf{A}^\delta)$ and then using Lemma .7(1), Corollary .9(2) and Lemma .11(2) respectively we see that $\varphi(k) \in \mathbb{K}(\mathbf{A}^\delta)$.

If $\psi(p)$ is of the form $\Box\psi_1(p)$, $\Diamond\psi_1(p)$ or $\blacksquare\psi_1(p)$ then $\psi_1(p)$ must be syntactically open and negative in p . We then use the inductive hypothesis to see that $\psi_1(k) \in \mathbb{O}(\mathbf{A}^\delta)$ and then use Corollary .9(1), Lemma .7(2), and Lemma .11(1) respectively to see that $\psi(k) \in \mathbb{O}(\mathbf{A}^\delta)$. \square

Lemma 8.2. *Let $\varphi(p, \bar{q}, \bar{\mathbf{i}}, \bar{\mathbf{m}})$ be syntactically closed and $\psi(p, \bar{q}, \bar{\mathbf{i}}, \bar{\mathbf{m}})$ be syntactically open. Let $D \subseteq \mathbb{K}(\mathbf{A}^\delta)$ be down-directed and let $U \subseteq \mathbb{O}(\mathbf{A}^\delta)$ be up-directed, let $\bar{b} \in A$, $\bar{c} \in J^\infty(\mathbf{A}^\delta)$ and $\bar{o} \in M^\infty(\mathbf{A}^\delta)$. Then*

1. (a) *if $\varphi(p, \bar{q}, \bar{\mathbf{i}}, \bar{\mathbf{m}})$ is positive in p , then $\varphi(\bigwedge D, \bar{b}, \bar{c}, \bar{o}) = \bigwedge \{ \varphi(d, \bar{b}, \bar{c}, \bar{o}) \mid d \in D \}$, and*
 (b) *if $\psi(p, \bar{q}, \bar{\mathbf{i}}, \bar{\mathbf{m}})$ is negative in p , then $\psi(\bigwedge D, \bar{b}, \bar{c}, \bar{o}) = \bigvee \{ \psi(d, \bar{b}, \bar{c}, \bar{o}) \mid d \in D \}$;*
2. (a) *if $\varphi(p, \bar{q}, \bar{\mathbf{i}}, \bar{\mathbf{m}})$ is negative in p , then $\varphi(\bigvee U, \bar{b}, \bar{c}, \bar{o}) = \bigwedge \{ \varphi(u, \bar{b}, \bar{c}, \bar{o}) \mid u \in U \}$, and*
 (b) *if $\psi(p, \bar{q}, \bar{\mathbf{i}}, \bar{\mathbf{m}})$ is positive in p , then $\psi(\bigvee U, \bar{b}, \bar{c}, \bar{o}) = \bigvee \{ \psi(u, \bar{b}, \bar{c}, \bar{o}) \mid u \in U \}$.*

Proof. We prove (1) by simultaneous induction on φ and ψ . As before, we will write $\varphi(p)$ for $\varphi(p, \bar{q}, \bar{\mathbf{i}}, \bar{\mathbf{m}})$ and $\psi(p)$ for $\psi(p, \bar{q}, \bar{\mathbf{i}}, \bar{\mathbf{m}})$.

The base cases of the induction for φ are when φ is of the form \top , \perp , p , a propositional variable q other than p , or \mathbf{i} . The base cases for ψ are those when ψ is of the form \top , \perp , a propositional variable q other than p , or \mathbf{m} . In each case the claim is trivially true.

If $\varphi(p) = \varphi_1(p) \vee \varphi_2(p)$ or $\varphi(p) = \varphi_1(p) \wedge \varphi_2(p)$ then φ_1 and φ_2 are syntactically closed and positive in p . Similarly, if $\psi(p) = \psi_1(p) \vee \psi_2(p)$ or $\psi(p) = \psi_1(p) \wedge \psi_2(p)$ then ψ_1 and ψ_2 are syntactically open and negative in p .

Thus when $\varphi(p) = \varphi_1(p) \wedge \varphi_2(p)$ the claim follows by the inductive hypothesis and the associativity of the meet operation, and when $\psi(p) = \psi_1(p) \vee \psi_2(p)$ the claim follows by the inductive hypothesis and the associativity of the join operation.

Now suppose that $\varphi(p) = \varphi_1(p) \vee \varphi_2(p)$. By the inductive hypothesis,

$$\varphi(\bigwedge D) = \varphi_1(\bigwedge D) \vee \varphi_2(\bigwedge D) = \left(\bigwedge_{d \in D} \varphi_1(d) \right) \vee \left(\bigwedge_{d \in D} \varphi_2(d) \right).$$

Since $\bigwedge \{ \varphi_1(d) \mid d \in D \} \leq \varphi_1(e)$ and $\bigwedge \{ \varphi_2(d) \mid d \in D \} \leq \varphi_2(e)$ for all $e \in D$ we have that $\varphi(\bigwedge D) \leq \bigwedge \{ \varphi_1(d) \vee \varphi_2(d) \mid d \in D \}$. For the reverse inequality, suppose that $x \in J^\infty(\mathbf{A}^\delta)$ and $x \not\leq \varphi(\bigwedge D)$. This implies that $x \not\leq \bigwedge \{ \varphi_1(d) \mid d \in D \}$ and $x \not\leq \bigwedge \{ \varphi_2(d) \mid d \in D \}$. Hence there exists $d_1, d_2 \in D$ such that $x \not\leq \varphi_1(d_1)$ and $x \not\leq \varphi_2(d_2)$. Since D is down-directed, there exists $d_3 \in D$ such that $d_3 \leq d_1$ and $d_3 \leq d_2$. Since φ_1 and φ_2 are monotone we have that $\varphi_1(d_3) \leq \varphi_1(d_1)$ and $\varphi_2(d_3) \leq \varphi_2(d_2)$. This implies that $x \not\leq \varphi_1(d_3)$ and $x \not\leq \varphi_2(d_3)$. By Lemma .3(3), $\varphi_1(d_3) \leq \kappa(x)$ and $\varphi_2(d_3) \leq \kappa(x)$ and so $\varphi_1(d_3) \vee \varphi_2(d_3) \leq \kappa(x)$. Hence $x \not\leq \varphi_1(d_3) \vee \varphi_2(d_3)$ and so $x \not\leq \bigwedge \{ \varphi_1(d) \vee \varphi_2(d) \mid d \in D \}$. Finally, by Lemma .4(1), $\bigwedge \{ \varphi(d) \mid d \in D \} \leq \varphi(\bigwedge D)$.

If $\psi(p) = \psi_1(p) \wedge \psi_2(p)$ then by the inductive hypothesis,

$$\psi(\bigwedge D) = \psi_1(\bigwedge D) \wedge \psi_2(\bigwedge D) = \left(\bigvee_{d \in D} \psi_1(d) \right) \wedge \left(\bigvee_{d \in D} \psi_2(d) \right).$$

Since $\bigvee \{ \psi_1(d) \mid d \in D \} \geq \psi_1(e)$ and $\bigvee \{ \psi_2(d) \mid d \in D \} \geq \psi_2(e)$ for all $e \in D$ we see that $\bigvee \{ \psi(d) \mid d \in D \} \leq \psi(\bigwedge D)$. For the reverse inequality, let $x \in J^\infty(\mathbf{A}^\delta)$ be such that $x \leq (\bigvee \{ \psi_1(d) \mid d \in D \}) \wedge (\bigvee \{ \psi_2(d) \mid d \in D \})$. Now $x \leq \bigvee \{ \psi_1(d) \mid d \in D \}$ and $x \leq \bigvee \{ \psi_2(d) \mid d \in D \}$ and so by Lemma .3(3) we have $\bigvee \{ \psi_1(d) \mid d \in D \} \not\leq \kappa(x)$ and $\bigvee \{ \psi_2(d) \mid d \in D \} \not\leq \kappa(x)$. This implies that there exist $d_1, d_2 \in D$ such that $\psi_1(d_1) \not\leq \kappa(x)$ and $\psi_2(d_2) \not\leq \kappa(x)$. Now by Lemma .3(3) we have $x \leq \psi_1(d_1)$ and $x \leq \psi_2(d_2)$. Furthermore, by the down-directedness of D , there exists $d_3 \in D$ such that $d_3 \leq d_1, d_2$. Since ψ_1 and ψ_2 are antitone in p we have that $x \leq \psi_1(d_3)$ and $x \leq \psi_2(d_3)$. This gives us that $x \leq \psi_1(d_3) \wedge \psi_2(d_3)$ and hence $x \leq \bigvee \{ \psi_1(d) \wedge \psi_2(d) \mid d \in D \}$.

Suppose $\varphi(p) = \mu^* X. \varphi_1(p, X)$ where $\varphi_1(p)$ is closed and positive in p . For each $d \in D$ we have $\{ a \in A \mid \varphi_1(d, a) \leq a \} \subseteq \{ a \in A \mid \varphi_1(\bigwedge D, a) \leq a \}$ and hence

$$\varphi(\bigwedge D) = \bigwedge \{ a \in A \mid \varphi_1(\bigwedge D, a) \leq a \} \leq \bigwedge_{d \in D} \left(\bigwedge \{ a \in A \mid \varphi_1(d, a) \leq a \} \right).$$

Now suppose that $\varphi_1(\wedge D, a) \leq a$. By the inductive hypothesis, $\varphi_1(\wedge D, a) = \wedge \{\varphi_1(d, a) \mid d \in D\}$ and so $\wedge \{\varphi_1(d, a) \mid d \in D\} \leq a$. By Lemma 8.1(1a), $\wedge \{\varphi_1(d, a) \mid d \in D\}$ is a meet of closed elements and hence closed. Thus we can apply compactness to get a finite set $\{d_i\}_{i=1}^n \subseteq D$ such that $\wedge_{i=1}^n \varphi_1(d_i, a) \leq a$. It follows that $\varphi_1(d_1 \wedge \dots \wedge d_n, a) \leq a$ and since D is down-directed, there exists $d_a \in D$ such that d_a is a lower bound for $\{d_i\}_{i=1}^n$. Now $\varphi_1(d_a, a) \leq a$.

Thus $\mu^* X. \varphi_1(d_a, X) = \wedge \{b \in A \mid \varphi_1(d_a, b) \leq b\} \leq a$ for each $a \in A$ with $\varphi_1(\wedge D, a) \leq a$. Finally this gives

$$\bigwedge_{d \in D} \varphi(d) = \bigwedge_{d \in D} \mu^* X. \varphi_1(d, X) \leq \bigwedge \{a \in A \mid \varphi_1(\wedge D, a) \leq a\} = \varphi(\wedge D).$$

If $\psi(p)$ is of the form $\nu^* X. \psi_1(p, X)$ then $\psi_1(p)$ must be open and negative in p . Now $\psi_1(d, a) \leq \psi_1(\wedge D, a)$ for all $d \in D$. Hence $\{a \in A \mid a \leq \psi_1(d, a)\} \subseteq \{a \in A \mid a \leq \psi_1(\wedge D, a)\}$ for all $d \in D$. This gives us

$$\bigvee_{d \in D} \psi(d) = \bigvee_{d \in D} \left\{ \bigvee \{a \in A \mid a \leq \psi_1(d, a)\} \right\} \leq \bigvee \{a \in A \mid a \leq \psi_1(\wedge D, a)\} = \psi(\wedge D).$$

Now suppose that $a \leq \psi_1(\wedge D, a)$. By the inductive hypothesis, we have that $\psi_1(\wedge D, a) = \vee \{\psi_1(d, a) \mid d \in D\}$, and by Lemma 8.1(1b) we have that $\psi_1(d, a)$ is open for each $d \in D$. Hence we can apply compactness to obtain a finite set $\{d_i\}_{i=1}^n \subseteq D$ such that $a \leq \vee_{i=1}^n \psi_1(d_i, a)$. Now since D is down-directed, there exists $d_a \in D$ such that d_a is a lower bound for the set $\{d_i \mid 1 \leq i \leq n\}$. Now $\psi_1(d_a, a) \geq \vee_{i=1}^n \psi_1(d_i, a)$ and hence $a \leq \psi_1(d_a, a)$. Now $a \leq \nu^* X. \psi_1(d_a, X)$ and so

$$\psi(\wedge D) = \bigvee \{a \in A \mid a \leq \psi_1(\wedge D, a)\} \leq \bigvee_{d \in D} \nu^* X. \psi_1(d, X) = \bigvee_{d \in D} \psi(d).$$

If $\varphi(p)$ is of the form $\odot \varphi_1(p)$ for $\odot \in \{\square, \diamond, \blacklozenge\}$, then $\varphi_1(p)$ must be syntactically closed and positive in p . By the inductive hypothesis, $\varphi_1(\wedge D) = \wedge \{\varphi_1(d) \mid d \in D\}$. By Lemma 8.1(1a) we have that $\varphi_1(d) \in \mathbb{K}(\mathbf{A}^\delta)$ for all $d \in D$. As φ_1 is monotone, and since D is down-directed, we have that $\{\varphi_1(d) \mid d \in D\}$ is a down-directed subset of $\mathbb{K}(\mathbf{A}^\delta)$. Thus we can use the fact that \square is \wedge -preserving, Lemma .8(2), and Lemma .10(2) to conclude that $\odot(\wedge \{\varphi_1(d) \mid d \in D\}) = \wedge \{\odot \varphi_1(d) \mid d \in D\}$ for $\odot \in \{\square, \diamond, \blacklozenge\}$.

If $\psi(p)$ is of the form $\odot \psi_1(p)$ for $\odot \in \{\square, \diamond, \blacksquare\}$, then $\psi_1(p)$ must be syntactically open and negative in p . By the inductive hypothesis, $\psi_1(\wedge D) = \vee_{d \in D} \psi_1(d)$. By Lemma 8.1(1b), each $\psi_1(d) \in \mathbb{O}(\mathbf{A}^\delta)$. Since D is down-directed, we have that $\{\psi_1(d) \mid d \in D\}$ is an up-directed set of open elements. We apply the fact that \square is \wedge -preserving, Lemma .8(2), and Lemma .10(2) to conclude that $\psi(\wedge D) = \odot \psi_1(\wedge D) = \odot(\vee \{\psi_1(d) \mid d \in D\}) = \vee \{\odot \psi_1(d) \mid d \in D\} = \vee \{\psi(d) \mid d \in D\}$.

If $\varphi(p)$ is of the form $\varphi_1(p) - \varphi_2(p)$ then $\varphi_1(p)$ is syntactically closed and positive in p , and $\varphi_2(p)$ is syntactically open and negative in p . By the inductive hypothesis we have $\varphi(\wedge D) = \varphi_1(\wedge D) - \varphi_2(\wedge D) = \wedge \{\varphi_1(d) \mid d \in D\} - \vee \{\varphi_2(e) \mid e \in D\}$. Now by Lemma 8.1(1a) and (1b), for all $d \in D$, we have $\varphi_1(d) \in \mathbb{K}(\mathbf{A}^\delta)$ and $\varphi_2(d) \in \mathbb{O}(\mathbf{A}^\delta)$. Since φ_1 is positive in p we have $\{\varphi_1(d) \mid d \in D\}$ is a down-directed set (of closed elements). Similarly, since φ_2 is negative in p we have $\{\varphi_2(e) \mid e \in D\}$ is an up-directed set (of open elements). Thus we can now apply Lemma .10(3) to get $\varphi(\wedge D) = \wedge \{\varphi_1(d) - \varphi_2(e) \mid d, e \in D\}$. Hence $\varphi(\wedge D) \leq \wedge \{\varphi_1(d) - \varphi_2(d) \mid d \in D\} = \wedge \{\varphi(d) \mid d \in D\}$. Now let $d, e \in D$. Since D is down-directed, there exists f such that $f \leq d$ and $f \leq e$. This gives us $\varphi_1(f) \leq \varphi_1(d)$ and $\varphi_2(e) \leq \varphi_2(f)$. By the fact that $-$ is order-preserving in its first argument and order-reversing in its second argument, we get that $\varphi_1(f) - \varphi_2(f) \leq \varphi_1(d) - \varphi_2(e)$. Thus we have $\varphi(\wedge D) = \wedge \{\varphi(d) \mid d \in D\}$.

If $\psi(p)$ is of the form $\psi_1(p) \rightarrow \psi_2(p)$ then $\psi_1(p)$ is syntactically closed and positive in p and $\psi_2(p)$ is syntactically open and negative in p . By the inductive hypothesis we have

$$\psi(\wedge D) = \psi_1(\wedge D) \rightarrow \psi_2(\wedge D) = \bigwedge \{\psi_1(d) \mid d \in D\} \rightarrow \bigvee \{\psi_2(e) \mid e \in D\}.$$

Again, $\{\psi_1(d) \mid d \in D\}$ is a down-directed subset of $\mathbb{K}(\mathbf{A}^\delta)$ and $\{\psi_2(e) \mid e \in D\}$ is an up-directed subset of $\mathbb{O}(\mathbf{A}^\delta)$. We apply Lemma .10(4) to get $\psi(\wedge D) = \bigvee \{\psi_1(d) \rightarrow \psi_2(e) \mid d, e \in D\}$. It is then clear that

$\psi(\wedge D) \geq \bigvee \{ \psi_1(d) \rightarrow \psi_2(d) \mid d \in D \}$. Since D is down-directed, for any $d, e \in D$, there exists a lower bound f for d, e . Now $\psi_1(f) \leq \psi_2(d)$ and $\psi_2(e) \leq \psi_2(f)$ and so $\psi_1(d) \rightarrow \psi_2(d) \leq \psi_1(f) \rightarrow \psi_2(f)$. Hence we conclude that $\psi(\wedge D) = \bigvee \{ \psi_1(f) \rightarrow \psi_2(f) \mid f \in D \} = \bigvee \{ \psi(d) \mid d \in D \}$. \square

We are now ready to prove the Ackermann lemmas which will justify the rules (RA) and (LA).

Lemma 8.3 (Righthanded Ackermann lemma for mu-algebras). *Let \mathbf{A} be a mu-algebra of the first kind. Let $\alpha(\bar{q}, \bar{\mathbf{i}}, \bar{\mathbf{n}})$, $\beta(p, \bar{q}, \bar{\mathbf{i}}, \bar{\mathbf{n}})$ and $\gamma(p, \bar{q}, \bar{\mathbf{i}}, \bar{\mathbf{n}})$ be \mathcal{L}_*^+ -formulas such that*

- (i) $\alpha(\bar{q}, \bar{\mathbf{i}}, \bar{\mathbf{n}})$ is syntactically closed and does not contain any occurrences of p ,
- (ii) $\beta(p, \bar{q}, \bar{\mathbf{i}}, \bar{\mathbf{n}})$ is syntactically closed and positive in p , and
- (iii) $\gamma(p, \bar{q}, \bar{\mathbf{i}}, \bar{\mathbf{n}})$ is syntactically open and negative in p .

Then for any $\bar{b} \in A$, $\bar{c} \in J^\infty(\mathbf{A}^\delta)$ and $\bar{d} \in M^\infty(\mathbf{A}^\delta)$, the following are equivalent

1. there exists $a \in A$ such that $\alpha(\bar{b}, \bar{c}, \bar{d}) \leq a$ and $\beta(a, \bar{b}, \bar{c}, \bar{d}) \leq \gamma(a, \bar{b}, \bar{c}, \bar{d})$,
2. $\beta(\alpha(\bar{b}, \bar{c}, \bar{d}), \bar{b}, \bar{c}, \bar{d}) \leq \gamma(\alpha(\bar{b}, \bar{c}, \bar{d}), \bar{b}, \bar{c}, \bar{d})$.

Proof. For the implication from top to bottom, it suffices to appeal to the monotonicity of β in p and the antitonicity of γ in p .

For the sake of the converse implication, assume that $\beta(\alpha(\bar{b}, \bar{c}, \bar{d}), \bar{b}, \bar{c}, \bar{d}) \leq \gamma(\alpha(\bar{b}, \bar{c}, \bar{d}), \bar{b}, \bar{c}, \bar{d})$. By Lemma 8.1, $\alpha(\bar{b}, \bar{c}, \bar{d})$ is closed. Hence, $\alpha(\bar{b}, \bar{c}, \bar{d}) = \bigwedge \{ u \in A \mid \alpha(\bar{b}, \bar{c}, \bar{d}) \leq u \}$, making it the meet of the down-directed family $U = \{ u \in A \mid \alpha(\bar{b}, \bar{c}, \bar{d}) \leq u \}$ of clopen elements. Thus we have

$$\beta(\bigwedge U, \bar{b}, \bar{c}, \bar{d}) \leq \gamma(\bigwedge U, \bar{b}, \bar{c}, \bar{d}).$$

Since β is syntactically closed and positive in p , and γ is syntactically open and negative in p , we may apply Lemma 8.2 to obtain

$$\bigwedge \{ \beta(u, \bar{b}, \bar{c}, \bar{d}) \mid u \in U \} \leq \bigvee \{ \gamma(u, \bar{b}, \bar{c}, \bar{d}) \mid u \in U \}.$$

By Lemma 8.1, $\beta(u, \bar{b}, \bar{c}, \bar{d})$ is closed and $\gamma(u, \bar{b}, \bar{c}, \bar{d})$ is open for each $u \in A$. Hence, by compactness,

$$\bigwedge_{i=1}^m \beta(u_i, \bar{b}, \bar{c}, \bar{d}) \leq \bigvee_{j=1}^n \gamma(u'_j, \bar{b}, \bar{c}, \bar{d}), \quad (7)$$

for some $u_1, \dots, u_m \in A$ with $\alpha(\bar{b}, \bar{c}, \bar{d}) \leq u_i$ for $1 \leq i \leq m$, and some $u'_1, \dots, u'_n \in A$ with $\alpha(\bar{b}, \bar{c}, \bar{d}) \leq u'_j$ for all $1 \leq j \leq n$.

Let $a = u_1 \wedge \dots \wedge u_m \wedge u'_1 \wedge \dots \wedge u'_n$. Then $\alpha(\bar{b}, \bar{c}, \bar{d}) \leq a \in A$. By the monotonicity of β in p , the antitonicity of γ in p , and (7), it follows that

$$\beta(a, \bar{b}, \bar{c}, \bar{d}) \leq \gamma(a, \bar{b}, \bar{c}, \bar{d}). \quad (8)$$

\square

Lemma 8.4 (Lefthanded Ackermann lemma for mu-algebras). *Let \mathbf{A} be a mu-algebra of the first kind. Let $\alpha(\bar{q}, \bar{\mathbf{i}}, \bar{\mathbf{n}})$, $\beta(p, \bar{q}, \bar{\mathbf{i}}, \bar{\mathbf{n}})$ and $\gamma(p, \bar{q}, \bar{\mathbf{i}}, \bar{\mathbf{n}})$ be \mathcal{L}^+ -formulas such that*

- (i) $\alpha(\bar{q}, \bar{\mathbf{i}}, \bar{\mathbf{n}})$ is syntactically open and does not contain any occurrences of p ,
- (ii) $\beta(p, \bar{q}, \bar{\mathbf{i}}, \bar{\mathbf{n}})$ is syntactically closed and negative in p , and
- (iii) $\gamma(p, \bar{q}, \bar{\mathbf{i}}, \bar{\mathbf{n}})$ is syntactically open and positive in p .

Then for any $\bar{b} \in A$, $\bar{c} \in J^\infty(\mathbf{A}^\delta)$ and $\bar{d} \in M^\infty(\mathbf{A}^\delta)$, the following are equivalent

1. there exists $a \in A$ such that $a \leq \alpha(\bar{b}, \bar{c}, \bar{d})$ and $\beta(a, \bar{b}, \bar{c}, \bar{d}) \leq \gamma(a, \bar{b}, \bar{c}, \bar{d})$,
2. $\beta(\alpha(\bar{b}, \bar{c}, \bar{d}), \bar{b}, \bar{c}, \bar{d}) \leq \gamma(\alpha(\bar{b}, \bar{c}, \bar{d}), \bar{b}, \bar{c}, \bar{d})$.

9 Canonicity of μ^* -ALBA inequalities

We now have all of the machinery in place to prove that admissible validity for an inequality $\varphi^* \leq \psi^*$ is preserved under any number of applications of rules from μ^* -ALBA. This will be crucial in showing both canonicity and tame canonicity for classes of inequalities on which the algorithm succeeds.

In this section we will assume that each of the inequalities that we are working with has already undergone preprocessing (see Section 5). The fact that both admissible validity and ordinary validity is preserved throughout the preprocessing stage is straightforward to prove.

Proposition 9.1. (Soundness of μ^* -ALBA rules w.r.t. admissible validity) *Let \mathbf{A} be a mu-algebra of the second kind, and $\varphi \leq \psi$ an \mathcal{L}_1 -inequality. Let $\text{FA}(\varphi^* \leq \psi^*) := \forall \mathbf{i} \forall \mathbf{m} (\mathbf{i} \leq \varphi^* \ \& \ \psi^* \leq \mathbf{m} \Rightarrow \mathbf{i} \leq \mathbf{m})$ and let $\{\text{Qlneq}_k \mid 1 \leq k \leq n\}$ be a set of quasi-inequalities obtained from $\text{FA}(\varphi^* \leq \psi^*)$ through the application of μ^* -ALBA rules. Then $\mathbf{A}^\delta \models_{\mathbf{A}} \varphi^* \leq \psi^*$ iff $\mathbf{A}^\delta \models_{\mathbf{A}} \{\text{Qlneq}_k \mid 1 \leq k \leq n\}$.*

Proof. We proceed by induction on the number of rule applications. For the base case, suppose that there are no applications of rules to $\text{FA}(\varphi^* \leq \psi^*)$. Assume $\mathbf{A}^\delta \models_{\mathbf{A}} \varphi^* \leq \psi^*$. Let V be an admissible assignment such that $V(\mathbf{i}) \leq V(\varphi^*)$ and $V(\psi^*) \leq V(\mathbf{m})$. By assumption we have that $V(\varphi^*) \leq V(\psi^*)$ and hence by transitivity we have $V(\mathbf{i}) \leq V(\mathbf{m})$.

Conversely, assume that $\mathbf{A}^\delta \models_{\mathbf{A}} \mathbf{i} \leq \varphi^* \ \& \ \psi^* \leq \mathbf{m} \Rightarrow \mathbf{i} \leq \mathbf{m}$. Let V be any admissible assignment. We need to show that $V(\varphi^*) \leq V(\psi^*)$. By the join-density of $J^\infty(\mathbf{A}^\delta)$ and the meet-density of $M^\infty(\mathbf{A}^\delta)$, it is sufficient to show that for every $i \in J^\infty(\mathbf{A}^\delta)$ and $m \in M^\infty(\mathbf{A}^\delta)$, if $i \leq V(\varphi^*)$ and $V(\psi^*) \leq m$, then $i \leq m$. Accordingly, suppose that $i \in J^\infty(\mathbf{A}^\delta)$ and $m \in M^\infty(\mathbf{A}^\delta)$ such that $i \leq V(\varphi^*)$ and $V(\psi^*) \leq m$. Let $V' \sim_{\mathbf{i}, \mathbf{m}} V$ such that $V'(\mathbf{i}) = i$ and $V'(\mathbf{m}) = m$. Then, $V'(\mathbf{i}) \leq V'(\varphi^*)$ and $V'(\psi^*) \leq V'(\mathbf{m})$, so by our assumption $i \leq m$.

Now suppose that $Q_1 = \{\text{Qlneq}_i \mid 1 \leq i \leq n\}$ is a set of quasi-inequalities obtained from $\text{FA}(\varphi^* \leq \psi^*)$ via k rule applications. Let $Q_2 = \{\text{Qlneq}_i \mid i \neq m, 1 \leq m \leq n\} \cup \{\text{Qlneq}_{m_j} \mid 1 \leq j \leq \ell\}$ be a set of quasi-inequalities obtained from Q_1 via a single rule application. (To be clear, Q_2 is the same as Q_1 except for the quasi-inequality $\text{Qlneq}_m \in Q_1$ which has been changed via the single rule application.)

We now cover the various possible cases for the $(k+1)$ -th rule application that gives us $\{\text{Qlneq}_{m_j} \mid 1 \leq j \leq \ell\}$ from Qlneq_m . The cases of the residuation and adjunction rules are clear from the properties of the residuals and adjoints.

We prove the case of $(\rightarrow \text{Appr})$. Let V be an admissible assignment such that $V(\chi \rightarrow \varphi) \leq V(\mathbf{m})$. We observe that $(\bigvee \{j \in J^\infty(\mathbf{A}^\delta) \mid j \leq V(\chi)\}) \rightarrow (\bigwedge \{n \in M^\infty(\mathbf{A}^\delta) \mid V(\varphi) \leq n\}) \leq V(\mathbf{m})$. Since \rightarrow is completely join-reversing in its first coordinate and completely meet-preserving in its second coordinate, we get $\bigwedge \{j \rightarrow n \mid j \leq V(\chi), V(\varphi) \leq n\} \leq V(\mathbf{m})$. As $V(\mathbf{m}) \in M^\infty(\mathbf{A}^\delta)$ and hence completely meet-prime, there must exist $j_0 \in J^\infty(\mathbf{A}^\delta)$ with $j_0 \leq V(\chi)$ and $n_0 \in M^\infty(\mathbf{A}^\delta)$ with $V(\varphi) \leq n_0$ such that $j_0 \rightarrow n_0 \leq V(\mathbf{m})$. Now choose $\mathbf{j} \in \text{NOM}, \mathbf{n} \in \text{CNOM}$ such that $\chi \rightarrow \varphi \leq \mathbf{m}$ is (\mathbf{j}, \mathbf{n}) -free and construct a (\mathbf{j}, \mathbf{n}) -variant V' of V such that $V'(\mathbf{j}) = j_0$ and $V'(\mathbf{n}) = n_0$. The opposite direction only requires the observation that \rightarrow is order-reversing in the first coordinate and order-preserving in the second coordinate.

The soundness and invertibility of the remaining ordinary approximation rules will follow using properties of canonical extension \mathbf{A}^δ , especially the join-density and join-primeness of $J^\infty(\mathbf{A}^\delta)$ and the meet-density and meet-primeness of $M^\infty(\mathbf{A}^\delta)$.

The case of the rule application being that of the Approximation Rules (μ^τ -A-R) and (ν^τ -A-R) follows from Proposition 7.3. We observe that these are the only rules which could possibly result in $\ell > 1$.

The case for the Ackermann rules (RA) and (LA) follow from Lemmas 8.3 and 8.4, respectively. \square

Corollary 9.2. *Let \mathbf{A} be a mu-algebra of the first kind, and $\varphi \leq \psi$ an \mathcal{L}_1 -inequality. Let Qlneq be a quasi-inequality obtained from $\text{FA}(\varphi^* \leq \psi^*)$ through the application of μ^* -ALBA rules none of which is (μ^τ -A-R) or (ν^τ -A-R). Then $\mathbf{A}^\delta \models_{\mathbf{A}} \varphi^* \leq \psi^*$ if and only if $\mathbf{A}^\delta \models_{\mathbf{A}} \text{Qlneq}$.*

Proof. The only part of the proof of Proposition 9.1 above that requires \mathbf{A} to be a mu-algebra of the second kind is the case of (μ^τ -A-R) and (ν^τ -A-R). Those cases rely on Proposition 7.3, which is proved using Lemma 7.2.

$$\begin{array}{ccc}
\mathbf{A} \models \varphi \leq \psi & & \mathbf{A}^\delta \models \varphi^* \leq \psi^* \\
\Downarrow & & \Uparrow \\
\mathbf{A}^\delta \models_{\mathbf{A}} \varphi^* \leq \psi^* & & \\
\Downarrow & & \\
\mathbf{A}^\delta \models_{\mathbf{A}} \text{pure}(\varphi^* \leq \psi^*) & \iff & \mathbf{A}^\delta \models \text{pure}(\varphi^* \leq \psi^*)
\end{array}$$

Figure 4: The proof of tame canonicity for \mathcal{L}_1 -inequalities on which a tame run of μ^* -ALBA succeeds.

The proof of that lemma requires \mathbf{A} be of the second kind. The soundness and invertibility of the remaining rules are proved exactly as in Proposition 9.1 and do not require \mathbf{A} to be of the second kind. \square

In order to complete the right-hand side of the U-shaped argument, we need to know that the rules can be ‘undone’ when using arbitrary assignments on the algebras that are the canonical extensions of mu-algebras of the first kind.

Proposition 9.3. *Let \mathbf{A} be a mu-algebra of the first kind. For any \mathcal{L}_1 -inequality $\varphi \leq \psi$ on which a tame run of μ^* -ALBA succeeds, we have that $\mathbf{A}^\delta \models \text{pure}(\varphi^* \leq \psi^*)$ if and only if $\mathbf{A}^\delta \models \varphi^* \leq \psi^*$.*

Proof. Since $\text{pure}(\varphi^* \leq \psi^*)$ is the result of a successful tame run of μ^* -ALBA, we know that $\text{pure}(\varphi^* \leq \psi^*)$ is a single quasi-inequality obtained via some finite number of rule applications from μ^* -ALBA, not including $(\mu^\tau\text{-A-R})$ or $(\nu^\tau\text{-A-R})$. Not including (RA) and (LA), the soundness and invertibility of each of the rules under arbitrary (i.e., not necessarily admissible) assignments is clear from the fact that the admissibility of the assignments was not used in any of the proofs in Proposition 9.1.

When working on perfect algebras, the Ackermann lemmas are easy to prove, and depend only on the monotonicity and antitonicity of the formulas involved. For (RA), the direction from top to bottom can be proved using only the fact that β is positive in p and γ is negative in p . The direction from bottom to top can be shown by taking V' to be the p -variant of V such that $V'(p) = V(\alpha)$. \square

Theorem 9.4. (Tame Canonicity) *All \mathcal{L}_1 -inequalities on which a tame run of μ^* -ALBA succeeds are tame canonical.*

Proof. Suppose that a tame run of μ^* -ALBA succeeds on the inequality $\varphi \leq \psi$. Then we have the following sequence of equivalences:

$$\begin{aligned}
& \mathbf{A} \models \varphi \leq \psi & (9) \\
\iff & \mathbf{A}^\delta \models_{\mathbf{A}} \varphi^* \leq \psi^* & (10) \\
\iff & \mathbf{A}^\delta \models_{\mathbf{A}} \text{pure}(\varphi^* \leq \psi^*) & (11) \\
\iff & \mathbf{A}^\delta \models \text{pure}(\varphi^* \leq \psi^*) & (12) \\
\iff & \mathbf{A}^\delta \models \varphi^* \leq \psi^* & (13)
\end{aligned}$$

The equivalence of (9) and (10) follows from the fact that the admissible validity of formulas without μ or ν on \mathbf{A}^δ will agree with ordinary admissible validity of formulas without μ or ν . For formulas with μ or ν , the definition of μ^* and ν^* assures that the interpretation of such formulas will be identical on \mathbf{A} and \mathbf{A}^δ . For the equivalence of (10) and (11) we use Corollary 9.2. The equivalence of (11) and (12) follows from the fact that $\text{pure}(\varphi^* \leq \psi^*)$ contains no propositional variables and hence that its admissible validity and validity coincide. The final equivalence is the statement of Proposition 9.3. \square

The next two lemmas and the proposition which follows are needed to convert formulas of the form $\mu^*X.\varphi(X)$ back to $\mu X.\varphi(X)$. The first of the lemmas is proved in [5].

Lemma 9.5. [5, Lemma 2.1] Let \mathbf{L} and \mathbf{M} be complete lattices and $G : M \times L \rightarrow L$. Let $\mu y.G(-, y) : \mathbf{M} \rightarrow \mathbf{L}$ be given for $a \in M$ by $a \mapsto \bigwedge \{x \in L \mid G(a, x) \leq x\}$.

If G is completely join-preserving, then $\mu y.G(-, y) : \mathbf{M} \rightarrow \mathbf{L}$ is defined everywhere on M and is completely join-preserving.

Lemma 9.6. Let \mathbf{A} be a mu-algebra of the second kind and let τ be an order type over n . Let $\psi(\bar{x}, X) \in \mathcal{L}_1$ such that $\psi(\bar{x}, X)$ is completely \vee -preserving in $(\bar{x}, X) \in (\mathbf{A}^\delta)^\tau \times \mathbf{A}^\delta$. If $\mathbf{A}^\delta \models [(\mathbf{i} \leq \mu^* X. \psi(\bar{\mathbf{j}}_i^\tau, X) \ \& \ \mathbf{j}^{\tau_i} \leq^{\tau_i} \bar{\varphi}_i) \Rightarrow \mathbf{i} \leq \mathbf{m}]$, then $\mathbf{A}^\delta \models [\mathbf{i} \leq \mu X. \psi(\bar{\varphi}, X) \Rightarrow \mathbf{i} \leq \mathbf{m}]$.

Proof. Assume that $\mathbf{A}^\delta \models [(\mathbf{i} \leq \mu^* X. \psi(\bar{\mathbf{j}}_i^\tau, X) \ \& \ \mathbf{j}^{\tau_i} \leq^{\tau_i} \bar{\varphi}_i) \Rightarrow \mathbf{i} \leq \mathbf{m}]$ and suppose $V(\mathbf{i}) \leq V(\mu X. \psi(\bar{\varphi}, X))$. By the join-density of $J^\infty(\mathbf{A}^\delta)$ we have that $\bar{\varphi} = \bigvee \{ \bar{j} \mid \bar{j} \in J^\infty((\mathbf{A}^\delta)^\tau), \bar{j} \leq \bar{\varphi} \}$. Thus $\mu X. \psi(\bar{\varphi}, X) = \mu X. \psi(\bigvee \{ \bar{j} \mid \bar{j} \leq \bar{\varphi} \}, X)$. Now by Lemma 9.5 with $M = (\mathbf{A}^\delta)^\tau$ and $L = \mathbf{A}^\delta$ we get $\mu X. \psi(\bar{\varphi}, X) = \bigvee \{ \mu X. \psi(\bar{j}, X) \mid \bar{j} \leq \bar{\varphi} \}$. Since $V(\mathbf{i}) \in J^\infty(\mathbf{A}^\delta)$, it is completely join-prime and so $V(\mathbf{i}) \leq \mu X. \psi(\bar{\varphi}, X) = \bigvee \{ \mu X. \psi(\bar{j}, X) \mid \bar{j} \leq \bar{\varphi} \}$ implies that there exists some \bar{j}_0 with $\bar{j}_0 \leq \bar{\varphi}$ such that $V(\mathbf{i}) \leq \mu X. \psi(\bar{j}_0, X) = \bigwedge \{ a \in \mathbf{A}^\delta \mid \psi(\bar{j}_0, a) \leq a \} \leq \bigwedge \{ a \in \mathbf{A} \mid \psi(\bar{j}_0, a) \} = \mu^* X. \psi(\bar{j}_0, X)$. Thus by the assumption, $V(\mathbf{i}) \leq V(\mathbf{m})$. \square

Proposition 9.7. Let \mathbf{A} be a mu-algebra of the second kind and $\varphi \leq \psi$ an \mathcal{L}_1 -inequality. Suppose that a proper run of μ^* -ALBA succeeds on $\varphi \leq \psi$, producing $\text{pure}(\varphi^* \leq \psi^*)$. If $\mathbf{A}^\delta \models \text{pure}(\varphi^* \leq \psi^*)$ then $\mathbf{A}^\delta \models \varphi \leq \psi$.

Proof. The set of quasi-inequalities $\text{pure}(\varphi^* \leq \psi^*)$ is obtained through a finite number of rule applications of μ^* -ALBA. As detailed by Proposition 9.3, all of the residuation, adjunction, ordinary approximation and Ackermann rules can be reversed while preserving validity on \mathbf{A}^δ . Suppose that at some stage during the proper run of μ^* -ALBA, the inequality $\mathbf{i} \leq \mu^* X. \psi(\bar{\varphi}/\bar{x}, X)$ is converted into $\bigvee_{i=1}^n (\exists \mathbf{j}^{\tau_i} [\mathbf{i} \leq \mu^* X. \psi(\bar{\mathbf{j}}_i^\tau/\bar{x}, X) \ \& \ \mathbf{j}^{\tau_i} \leq^{\tau_i} \varphi_i])$. By Lemma 9.6, the disjunction $\bigvee_{i=1}^n (\exists \mathbf{j}^{\tau_i} [\mathbf{i} \leq \mu^* X. \psi(\bar{\mathbf{j}}_i^\tau/\bar{x}, X) \ \& \ \mathbf{j}^{\tau_i} \leq^{\tau_i} \varphi_i])$ can be converted into $\mathbf{i} \leq \mu X. \psi(\bar{\varphi}, X)$. \square

We are now ready to present the final canonicity result.

Theorem 9.8. (Canonicity) Let \mathbf{A} be a mu-algebra of the second kind and let $\varphi \leq \psi$ be an \mathcal{L}_1 -inequality on which a proper run of μ^* -ALBA succeeds. If $\mathbf{A} \models \varphi \leq \psi$ then $\mathbf{A}^\delta \models \varphi \leq \psi$.

Proof. The proof is similar to that of Theorem 9.4.

$$\mathbf{A} \models \varphi \leq \psi \tag{14}$$

$$\iff \mathbf{A}^\delta \models_{\mathbf{A}} \varphi^* \leq \psi^* \tag{15}$$

$$\iff \mathbf{A}^\delta \models_{\mathbf{A}} \text{pure}(\varphi^* \leq \psi^*) \tag{16}$$

$$\iff \mathbf{A}^\delta \models \text{pure}(\varphi^* \leq \psi^*) \tag{17}$$

$$\implies \mathbf{A}^\delta \models \varphi \leq \psi \tag{18}$$

The equivalence of (14) and (15) follows as for Theorem 9.4. For the equivalence of (15) and (16) we apply the inductive argument from Proposition 9.1. The equivalence of (16) and (17) follows from the fact that $\text{pure}(\varphi^* \leq \psi^*)$ contains no propositional variables and hence admissible validity and validity coincide for it. The final implication is the statement of Proposition 9.7. \square

Lastly, we make an observation regarding mu-inequalities with no occurrences of fixed point binders. If $\varphi \leq \psi$ does not contain any fixed point binders, and a tame run of μ^* -ALBA succeeds on $\varphi \leq \psi$, then $\varphi \leq \psi$ will be tame canonical. That is, $\mathbf{A} \models \varphi \leq \psi$ if and only if $\mathbf{A}^\delta \models \varphi^* \leq \psi^*$. However, we have that $\varphi = \varphi^*$ and $\psi = \psi^*$ and so $\mathbf{A} \models \varphi \leq \psi$ if and only if $\mathbf{A}^\delta \models \varphi \leq \psi$. Likewise, if a proper run of μ^* -ALBA succeeds on $\alpha \leq \beta$ where this mu-inequality has no occurrences of fixed point binders, Proposition 9.3 can be applied to show the converse of (21) \Rightarrow (22) in Theorem 9.8 and hence $\mathbf{A} \models \alpha \leq \beta$ if and only if $\mathbf{A}^\delta \models \alpha \leq \beta$.

10 Canonicity of the restricted and tame inductive mu-inequalities

In this section we argue that μ^* -ALBA successfully purifies all restricted inductive mu-inequalities by means of proper runs, and all tame inductive mu-inequalities by means of tame runs. Once we have established these claims, the next theorem will then follow from Theorems 9.4 and 9.8:

Theorem 10.1. *All restricted inductive mu-inequalities are canonical and all tame inductive mu-inequalities are tame canonical.*

As far as tame inductive inequalities are concerned, the proof is almost verbatim the same as in [6, Section 10.1]. Indeed, tame inductive inequalities are just (a subset of the) inductive inequalities from the language of intuitionistic modal logic with some fixed point binders ‘along for the ride’ on the non-critical branches. Thus the binders never need to be handled by the algorithm, so the runs are guaranteed to be tame and exactly the same strategy employed in [6] is sufficient. The only point that needs to be checked is that the side conditions of the Ackermann rules pertaining to syntactic openness and closure are met when these rules need to be applied. This follows from the next two lemmas:

Lemma 10.2. *If μ^* -ALBA is applied to any \mathcal{L}_1 -inequality then, during the whole run, every inequality in every antecedent of every produced quasi-inequality is either pure, or has a syntactically almost closed left-hand side and a syntactically almost open right-hand side. Consequently, if a non-pure inequality contains no fixed point binders it has a syntactically closed left-hand side and a syntactically open right-hand side.*

Proof. The proof is by induction on the application of the rules of μ^* -ALBA. Suppose μ^* -ALBA is run on an \mathcal{L}_1 -inequality. As discussed above, preprocessing turns this into a finite number of inequalities $\eta \leq \beta$. Starring and first approximation produce $(\mathbf{i} \leq \eta^* \ \& \ \beta^* \leq \mathbf{m}) \Rightarrow \mathbf{i} \leq \mathbf{m}$. We claim that $\mathbf{i} \leq \eta^*$ and $\beta^* \leq \mathbf{m}$ satisfy the statement of the lemma. Indeed, \mathbf{i} and \mathbf{m} are, respectively, syntactically closed and open. As $\eta^*, \beta^* \in \mathcal{L}_*$, they do not contain any nominals, co-nominals, \blacklozenge or \blacksquare , and are therefore both syntactically almost open and syntactically almost closed.

The induction now proceeds by showing that the desired properties are invariant under the application of the rules of μ^* -ALBA. The most interesting cases are those for $(\mu^t\text{-A-R})$ and $(\nu^t\text{-A-R})$. We verify $(\mu^t\text{-A-R})$. Applied to an inequality $\mathbf{i} \leq \mu^*X.\psi(\overline{\varphi}/\overline{x}, X, \overline{y}/\overline{z})$, the rule $(\mu^t\text{-A-R})$ produces a disjunction of inequalities of the form $\mathbf{i} \leq \mu^*X.\psi(\overline{\mathbf{j}}_i/\overline{x}, X, \overline{y}/\overline{z})$ and $\mathbf{j}^{r_i} \leq \varphi_i$. Since the \overline{y} are constant sentences, the first of these inequalities is pure. If $r_i = 1$, the second inequality is $\mathbf{j} \leq \varphi_i$ where \mathbf{j} is syntactically closed and φ_i is syntactically almost open since it occurs positively as a subformula of the syntactically almost open formula $\mu^*X.\psi(\overline{\varphi}/\overline{x}, X, \overline{y}/\overline{z})$. On the other hand, if $r_i = \partial$, the second inequality is $\varphi_i \leq \mathbf{n}$ where \mathbf{n} is syntactically open and φ_i is syntactically almost closed since it occurs negatively as a subformula of the syntactically almost open formula $\mu^*X.\psi(\overline{\varphi}/\overline{x}, X, \overline{y}/\overline{z})$. \square

Lemma 10.3. *If μ^* -ALBA is applied to any tame inductive \mathcal{L}_1 inequality, then, during the whole run, every inequality in every antecedent of every quasi-inequality is either pure or has a left-hand side (resp., right-hand side) in which all occurrences of μ^* are positive (resp., negative) and all occurrences of ν^* are negative (resp., positive).*

Proof. Preprocessing turns any tame inductive \mathcal{L}_1 inequality into a finite number of tame inductive inequalities $\eta \leq \beta$. By the definition of tame inductive inequalities, the only occurrences of binder nodes in $+\eta$ and $-\beta$ are $+\nu$ and $-\mu$. Thus starring and first approximation yields a quasi-inequality

$$(\mathbf{i} \leq \eta^* \ \& \ \beta^* \leq \mathbf{m}) \Rightarrow \mathbf{i} \leq \mathbf{m},$$

where in η^* (resp., β^*) all occurrences of all occurrences of μ^* are negative (resp., positive) and all occurrences of ν^* are positive (resp., negative). Thus this quasi-inequality satisfies the claim of the lemma. It is now sufficient to show that these conditions are invariant under the application of all μ^* -ALBA rules. For the sake of the adjunction rules, note that all occurrences of binders are untouched and left on the same sides of inequalities. The residuation rules move subformulas across the inequality but with an accompanying change in polarity, and

the same consideration also deals with the Ackermann rules. The fixed point approximation rules (μ^r -A-R) and (ν^r -A-R) are not applied, since tame inductive inequalities contain no binders on critical branches. The cases for the other approximation rules are easy to verify — let us look explicitly only at (\rightarrow Appr). By assumption, in the premise $\chi \rightarrow \varphi \leq \mathbf{m}$, all occurrences of μ^* are positive and all occurrences of ν^* are negative. Thus, in $\mathbf{j} \rightarrow \mathbf{n} \leq \mathbf{m} \ \& \ \mathbf{j} \leq \chi \ \& \ \varphi \leq \mathbf{m}$, the conclusion of the rule, the inequality $\mathbf{j} \rightarrow \mathbf{n}$ is pure, the righthand side of $\mathbf{j} \leq \chi$ contains μ^* only negatively and ν^* only positively, while the lefthand side of $\varphi \leq \mathbf{m}$ contains μ^* only positively and ν^* only negatively. \square

Now that we have established Theorem 10.1 for tame inductive mu-inequalities, we will focus on the restricted inductive mu-inequalities for the remainder of this section. In [5, Section 9] it was proven that μ -ALBA successfully purifies all recursive mu-inequalities. The restricted inductive mu-inequalities form a subclass of these, but μ^* -ALBA is also a restricted version of μ -ALBA. Consequently our argument will follow that in [5] very closely, but we will be at pains to show how the constraints built into the restricted inductive mu-inequalities allow us to still succeed with the restricted resources of μ^* -ALBA. The reader might find it useful to refer to Example 11.1 while reading the following argument.

Let $\eta \leq \beta$ be an (Ω, ϵ) -inductive inequality. We proceed as in ALBA and preprocess this inequality by applying splitting and (\top) and (\perp) exhaustively. This might produce multiple inequalities, on each of which we proceed separately. On each such inequality, denoted again $\eta \leq \beta$, we proceed to first approximation, which yields the following quasi-inequality:

$$\forall \bar{p} \forall \mathbf{i} \forall \mathbf{m} [(\mathbf{i} \leq \eta \ \& \ \beta \leq \mathbf{m}) \Rightarrow \mathbf{i} \leq \mathbf{m}]. \quad (19)$$

Because its consequent is always pure, we only concentrate on its antecedent. Since the outer skeletons of β and η are built exactly as the outer part of an inductive modal formula, the ordinary approximation rules can be applied so as to surface the inner skeleton. So we can equivalently rewrite $\mathbf{i} \leq \eta \ \& \ \beta \leq \mathbf{m}$ as the conjunction of a set of inequalities which, whenever they contain critical variables in the scope of fixed point binders occurring as skeleton nodes, are of the form

$$\mathbf{i} \leq \mu X. \psi'(\bar{p}) \quad \text{and} \quad \nu X. \varphi'(\bar{p}) \leq \mathbf{m}, \quad (20)$$

where $\mu X. \psi'(\bar{p})$ and $\nu X. \varphi'(\bar{p})$ are sentences. (For the critical branches which do not contain such fixed point binders, we further proceed by exhaustively applying the approximation rules as in ALBA in order to surface the PIA parts.)

In what follows, we call a generation tree *non-trivially restricted* (Ω, ϵ) -inductive if it is restricted (Ω, ϵ) -inductive and contains at least one ϵ -critical branch.

Proposition 10.4. 1. The inequality $\mathbf{i} \leq \mu X. \psi'$ in (20) is of the form $\mathbf{i} \leq \mu X. \psi(\bar{\varphi}/\bar{y}, X, \bar{\gamma}/\bar{z})$, where $\mu X. \psi(\bar{y}, X, \bar{z})$ is an (\bar{y}, \bar{z}) -IF $_{\tau}^{\diamond}$ formula for some order-type τ over \bar{y} , and the $\bar{\gamma}$ are constant sentences;
2. the inequality $\nu X. \varphi' \leq \mathbf{m}$ in (20) is of the form $\nu X. \varphi(\bar{\psi}/\bar{y}, X, \bar{\gamma}/\bar{z}) \leq \mathbf{m}$, where $\nu X. \varphi(\bar{y}, X, \bar{z})$ is an (\bar{y}, \bar{z}) -IF $_{\tau}^{\square}$ formula for some order-type τ over \bar{y} , and the $\bar{\gamma}$ are constant sentences.

Proof. Notice that preprocessing, first approximation and ordinary approximation rules do not involve fixed points. Hence a proof very similar to that of [6, Lemma 10.6] proves that $+\mu X. \psi'$ and $-\nu X. \varphi'$ are non-trivially restricted (Ω, ϵ) -inductive. Hence the statement immediately follows from Lemma 10.5 below. \square

Lemma 10.5. 1. Let ψ' be such that $+\psi'$ is non-trivially restricted (Ω, ϵ) -inductive, and the P_3 -paths of all critical branches are of length 0. Then ψ' is of the form $\psi(\bar{\varphi}/\bar{y}, \bar{X}, \bar{\gamma}/\bar{z})$ where $\psi(\bar{x}, \bar{z})$ is an (\bar{x}, \bar{z}) -IF $_{\tau}^{\diamond}$ formula, for $\bar{x} = \bar{y} \oplus \bar{X}$ and some order-type τ over \bar{x} , the $\bar{\varphi}$ are sentences and the $\bar{\gamma}$ are constant sentences. Moreover, if $\bar{y} = (y_1, \dots, y_n)$ then, for each $1 \leq i \leq n$, $+\varphi_i$ is ϵ -PIA if $\tau_i = 1$ and $-\varphi_i$ is ϵ -PIA if $\tau_i = \emptyset$. Finally $\epsilon^{\partial}(+\psi(\bar{x}, \bar{\gamma}/\bar{z}))$.

2. Let φ' be such that $-\varphi'$ is non-trivially restricted (Ω, ϵ) -inductive, and the P_3 -paths of all critical branches are of length 0. Then φ' is of the form $\varphi(\bar{\psi}/\bar{y}, \bar{X}, \bar{y}/\bar{z})$ where $\varphi(\bar{x}, \bar{z})$ is an (\bar{x}, \bar{z}) -IF $_{\tau}^{\diamond}$ formula, for $\bar{x} = \bar{y} \oplus \bar{X}$ and some order-type τ over \bar{x} , the $\bar{\varphi}$ are sentences and the \bar{y} are constant sentences. Moreover, if $\bar{y} = (y_1, \dots, y_n)$ then, for each $1 \leq i \leq n$, $-\psi_i$ is ϵ -PIA if $\tau_i = 1$ and $+\psi_i$ is ϵ -PIA if $\tau_i = \partial$. Finally $\epsilon^{\partial}(-\varphi(\bar{x}, \bar{y}/\bar{z}))$.

Proof. The proof is virtually identical to that of [5, Lemma 9.2] and [5, Remark 9.3], the only difference being that we need to ensure that the \bar{y} are *constant*, i.e., contain no propositional variables. We will do so by appealing to the condition (NL) at the appropriate point. The proof is by simultaneous induction on the *skeleton depths* of $+\psi'$ and $-\varphi'$, where the skeleton depth of an ϵ -recursive (or inductive) generation tree $*\xi$, with $*$ $\in \{+, -\}$, is the maximum length of the P_2 parts of ϵ -critical branches in $*\xi$.

Consider the base case when the skeleton depth of $+\psi'$ is 0, i.e., the critical branches consist only of PIA nodes. Then, by (GB1), ψ' is a sentence. Hence we let $\psi = x_1$ which is IF $_{\tau}^{\diamond}$ with $\tau = (1)$, $\varphi_1 = \psi'$ (which is ϵ -PIA and a sentence). Moreover, (vacuously) all \bar{y} are constant sentences.

We assume that the claim holds for all $+\psi'$ and $-\varphi'$ of skeleton depth at most k . The inductive step consists of checking the possible cases (corresponding to the possible main connectives of ψ' and φ'). We verify the case when $-\varphi'$ is of the form $-(\psi' \rightarrow \chi)$. By (GB3) and (NL), χ is a constant sentence (and $\epsilon^{\partial}(-\chi)$). Then $+\psi'$ is non-trivially restricted (Ω, ϵ) -inductive, and hence, by the induction hypothesis, ψ' is of the form $\psi(\bar{\varphi}/\bar{y}, \bar{X}, \bar{y}/\bar{z})$, where $\psi(\bar{x}, \bar{z})$ is an (\bar{x}, \bar{z}) -IF $_{\tau}^{\diamond}$ formula for some order-type τ over $\bar{x} = \bar{y} \oplus \bar{X}$, and the $\bar{\varphi}$ are sentences and the \bar{y} are constant sentences. Moreover, $\bar{y} = (y_1, \dots, y_n)$ and for every $1 \leq i \leq n$, the generation tree $-\varphi_i$ is non-trivially (Ω, ϵ) -PIA if $(\tau^{\partial})_i = 1$ (i.e., $\tau_i = \partial$) and $+\varphi_i$ is non-trivially (Ω, ϵ) -PIA if $\tau_i^{\partial} = \partial$ (i.e., $\tau_i = 1$). Then we let $\varphi = \psi(\bar{x}, \bar{z}) \rightarrow z$, which is $(\bar{x}, \bar{z} \oplus z)$ -IF $_{\tau}^{\diamond}$, where z is a fresh variable. Moreover, for $1 \leq i \leq n$ we let $\psi_i = \psi_i$. Hence φ' is of the form $(\psi(\bar{\psi}/\bar{y}, \bar{X}, \bar{y}/\bar{z}) \rightarrow z)[\chi/z]$, with $\varphi_1 \dots \varphi_n$ (obtained above from the induction hypothesis) playing the role of $\psi_1 \dots \psi_n$. Finally, $\epsilon^{\partial}(\psi(\bar{x}, \bar{y}/\bar{z}) \rightarrow \chi)$, since $\epsilon^{\partial}(-\chi)$, and the induction hypothesis implies that $\epsilon^{\partial}(+\psi(\bar{x}, \bar{y}/\bar{z}))$. \square

Proposition 10.4 implies that the approximation rules $(\mu^{\tau}\text{-A-R})$ and $(\nu^{\tau}\text{-A-R})$ can be applied to the inequalities (20), respectively.² In addition to this, we can assume w.l.o.g. that every inequality sitting in an antecedents of any quasi-inequality produced by these rule applications and containing a critical branch is of the form

$$\mathbf{j} \leq \varphi \quad \text{or} \quad \psi \leq \mathbf{n}, \quad (21)$$

where $+\varphi$ and $-\psi$ are non-trivially (Ω, ϵ) -PIA (i.e., non-trivially (Ω, ϵ) -inductive with all critical branches consisting only of PIA-nodes, i.e., of P_1 -nodes) and, by (NB-PIA) and (GB1) φ and ψ are sentences containing no binders, i.e., they are formulas of \mathcal{L} . In other words φ and ψ are (Ω, ϵ) -inductive formulas of intuitionistic modal logic. This means that we can proceed as in ALBA to eliminate all propositional variables from the quasi-inequality. However, before the Ackermann rules can be applied we need to be sure that their side conditions are met. This is guaranteed by the following considerations: after the application of the approximation rules $(\mu^{\tau}\text{-A-R})$ and $(\nu^{\tau}\text{-A-R})$, all fixed point binders in the current quasi-inequalities occur only in pure inequalities of the form $\mathbf{i} \leq \mu^* X. \psi(\bar{\mathbf{j}}_i^{\epsilon}/\bar{x}, X, \bar{y}/\bar{z})$ and $\nu^* X. \varphi(\bar{\mathbf{n}}_i^{\epsilon}/\bar{x}, X, \bar{y}/\bar{z}) \leq \mathbf{m}$. The only inequalities involved in the application of an Ackermann rule are non-pure, and therefore cannot contain any fixed point binders. So by Lemma 10.2 above, their left hand sides will be syntactically closed and their right and sides syntactically open, as desired.

In summary: In dealing with the outer skeleton we proceed as in ALBA. The inner skeleton is processed as in μ -ALBA, but with the applicability of the restricted approximation rules $(\mu^{\tau}\text{-A-R})$ and $(\nu^{\tau}\text{-A-R})$ guaranteed by Proposition 10.4 via Lemma 10.5. This proposition and lemma are provable thanks to the condition (NL). The condition (NB-PIA) ensures that the PIA parts contain no fixed point binders or fixed point variables and

²Applying one of these approximation rules within the antecedent of a quasi-inequality may split that quasi-inequality into the conjunction of several quasi-inequalities, on each of which we proceed separately.

can hence be treated as in ALBA, without the need of the adjunction rules for μ and ν used in μ -ALBA. Next, the condition (Ω -CONF) guarantees that we are working with inductive rather than recursive inequalities, and hence that the ordinary (non-recursive) Ackermann rule can be used to eliminate the propositional variables. Lastly, the requirement that every occurrence of a binder is on an ϵ -critical branch (Definition 4.6) guarantees that every binder must have one of the approximation rules (μ^r -A-R) and (ν^r -A-R) applied to it, i.e., that the run of the algorithm is proper.

11 Examples

Example 11.1. Let us consider the restricted inductive inequality $\diamond \mu X.(\diamond X \vee \square(\square \diamond q \vee p)) \leq \nu Y.([\square((q \rightarrow \perp) \wedge (p \rightarrow \perp)) \rightarrow \perp] \wedge \square Y)$ from Example 4.8. No preprocessing is possible, so starring the binders and applying the first approximation rule produces:

$$\forall \mathbf{i} \forall \mathbf{m} [\mathbf{i} \leq \diamond \mu^* X.(\diamond X \vee \square(\square \diamond q \vee p)) \& \nu^* Y.([\square((q \rightarrow \perp) \wedge (p \rightarrow \perp)) \rightarrow \perp] \wedge \square Y) \leq \mathbf{m} \Rightarrow \mathbf{i} \leq \mathbf{m}] \quad (22)$$

We begin by applying approximation rules to surface the PIA parts. The (\diamond Appr) rule transforms the first inequality in the antecedent into $\mathbf{i} \leq \diamond \mathbf{j}$ and $\mathbf{j} \leq \mu^* X.(\diamond X \vee \square(\square \diamond q \vee p))$. We have to apply (μ^r -A-R) to the latter inequality. To that end, consider the subformula $(\diamond X \vee \square(\square \diamond q \vee p))$ — it is of the form $\psi(X/x_1, \square(\square \diamond q \vee p)/x_2)$ where $\psi(x_1, x_2) = \diamond x_1 \vee x_2$, which is completely \vee -preserving as a map from $\mathbf{C} \times \mathbf{C}$ to \mathbf{C} (indeed, $\psi(x_1, x_2)$ is an (x_1, x_2) -IF $_{(1,1)}^\diamond$ formula). So we may apply (μ^r -A-R) to produce $\mathbf{j} \leq \mu^* X.(\diamond X \vee \mathbf{k})$ and $\mathbf{k} \leq \square(\square \diamond q \vee p)$.

Next we have to apply (ν^r -A-R) to $\nu^* Y.([\square((q \rightarrow \perp) \wedge (p \rightarrow \perp)) \rightarrow \perp] \wedge \square Y) \leq \mathbf{m}$. The subformula $[\square((q \rightarrow \perp) \wedge (p \rightarrow \perp)) \rightarrow \perp] \wedge \square Y$ is of the form $\varphi(\square((q \rightarrow \perp) \wedge (p \rightarrow \perp))/x_1, Y/x_2)$ where $\varphi(x_1, x_2) = (x_1 \rightarrow \perp) \wedge \square x_2$. Now $\varphi(x_1, x_2)$ is completely \wedge -preserving as a map from $\mathbf{C}^\partial \times \mathbf{C}$ to \mathbf{C} — one way to see this is to note that it is an (x_1, x_2) -IF $_{(\partial,1)}^\square$ formula. Applying (ν^r -A-R) yields $\nu^* Y.([\mathbf{l} \rightarrow \perp] \wedge \square Y) \leq \mathbf{m}$ and $\mathbf{l} \leq \square((q \rightarrow \perp) \wedge (p \rightarrow \perp))$.

Thus the PIA parts have been surfaced and the quasi-inequality has been transformed into:

$$\left[\& \left(\begin{array}{ll} \mathbf{i} \leq \diamond \mathbf{j} & \mathbf{k} \leq \square(\square \diamond q \vee p) \\ \mathbf{j} \leq \mu^* X.(\diamond X \vee \mathbf{k}) & \mathbf{l} \leq \square((q \rightarrow \perp) \wedge (p \rightarrow \perp)) \\ \nu^* Y.([\mathbf{l} \rightarrow \perp] \wedge \square Y) \leq \mathbf{m} & \end{array} \right) \Rightarrow \mathbf{i} \leq \mathbf{m} \right].$$

We now proceed to apply adjunction and residuation rules to get the quasi-inequality into the right shape for the application of the Ackermann rules. Applying (\square RA), (\vee RR) and (\wedge RA) produces

$$\left[\& \left(\begin{array}{lll} \mathbf{i} \leq \diamond \mathbf{j} & \blacklozenge \mathbf{k} - \square \diamond q \leq p & \\ \mathbf{j} \leq \mu^* X.(\diamond X \vee \mathbf{k}) & \blacklozenge \mathbf{l} \leq q \rightarrow \perp & \blacklozenge \mathbf{l} \leq p \rightarrow \perp \\ \nu^* Y.([\mathbf{l} \rightarrow \perp] \wedge \square Y) \leq \mathbf{m} & & \end{array} \right) \Rightarrow \mathbf{i} \leq \mathbf{m} \right].$$

Now applying (\rightarrow RA) and then (\wedge LR) to $\blacklozenge \mathbf{l} \leq q \rightarrow \perp$ transforms it into $q \leq \blacklozenge \mathbf{l} \rightarrow \perp$. This yields the following quasi-inequality, which has now been solved for $+p$ and $-q$ and is therefore ready for the application of the right and lefthanded Ackermann rules (RA) and (LA) to eliminate p and q , respectively:

$$\left[\& \left(\begin{array}{lll} \mathbf{i} \leq \diamond \mathbf{j} & \blacklozenge \mathbf{k} - \square \diamond (\blacklozenge \mathbf{l} \rightarrow \perp) \leq p & \\ \mathbf{j} \leq \mu^* X.(\diamond X \vee \mathbf{k}) & q \leq \blacklozenge \mathbf{l} \rightarrow \perp & \blacklozenge \mathbf{l} \leq p \rightarrow \perp \\ \nu^* Y.([\mathbf{l} \rightarrow \perp] \wedge \square Y) \leq \mathbf{m} & & \end{array} \right) \Rightarrow \mathbf{i} \leq \mathbf{m} \right].$$

Now applying (LA) gives

$$\left[\& \left(\begin{array}{ll} \mathbf{i} \leq \diamond \mathbf{j} & \blacklozenge \mathbf{k} - \square \diamond (\blacklozenge \mathbf{l} \rightarrow \perp) \leq p \\ \mathbf{j} \leq \mu^* X.(\diamond X \vee \mathbf{k}) & \blacklozenge \mathbf{l} \leq p \rightarrow \perp \\ \nu^* Y.([\mathbf{l} \rightarrow \perp] \wedge \square Y) \leq \mathbf{m} & \end{array} \right) \Rightarrow \mathbf{i} \leq \mathbf{m} \right],$$

and then applying (RA) gives the pure quasi-inequality

$$\left[\& \left(\begin{array}{l} \mathbf{i} \leq \Diamond \mathbf{j} \\ \mathbf{j} \leq \mu^* X.(\Diamond X \vee \mathbf{k}) \\ \nu^* Y.(\mathbf{l} \rightarrow \perp) \wedge \Box Y \leq \mathbf{m} \\ \blacklozenge \mathbf{l} \leq (\blacklozenge \mathbf{k} - \Box \Diamond(\blacklozenge \mathbf{l} \rightarrow \perp)) \rightarrow \perp \end{array} \right) \Rightarrow \mathbf{i} \leq \mathbf{m} \right].$$

Example 11.2. Consider the tame inductive inequality $\Diamond(\Box \perp \vee p) \wedge \Box q \leq \mu Y.(\Diamond(p \wedge q) \wedge \Box Y)$ from Example 4.10. Preprocessing distributes \Diamond and \wedge over \vee in the antecedent producing $(\Diamond \Box \perp \wedge \Box q) \vee (\Diamond p \wedge \Box q) \leq \mu Y.(\Diamond(p \wedge q) \wedge \Box Y)$, to which we can apply (\vee LA) to split into two inequalities: $(\Diamond \Box \perp \wedge \Box q) \leq \mu Y.(\Diamond(p \wedge q) \wedge \Box Y)$ and $(\Diamond p \wedge \Box q) \leq \mu Y.(\Diamond(p \wedge q) \wedge \Box Y)$. As p appears only positively in the first, we may eliminate it by applying the (\perp) rule to obtain $(\Diamond \Box \perp \wedge \Box q) \leq \mu Y.(\Diamond(\perp \wedge q) \wedge \Box Y)$. The algorithm μ^* -ALBA now proceeds separately on these two inequalities. We will describe this process only for the second one, the execution on the first being very similar but less interesting. Starring and first approximation transforms the second inequality into the quasi-inequality

$$\mathbf{i} \leq \Diamond p \wedge \Box q \ \& \ \mu^* Y.(\Diamond(p \wedge q) \wedge \Box Y) \leq \mathbf{m} \Rightarrow \mathbf{i} \leq \mathbf{m}.$$

Applying (\wedge RA) to $\mathbf{i} \leq \Diamond p \wedge \Box q$ produces $\mathbf{i} \leq \Diamond p$ and $\mathbf{i} \leq \Box q$. To the first of these we apply (\Diamond Appr) to produce $\mathbf{i} \leq \Diamond \mathbf{j}$ and $\mathbf{j} \leq p$, while (\Box RA) turns the second into $\blacklozenge \mathbf{i} \leq q$. Thus we have the quasi-inequality

$$\mathbf{i} \leq \Diamond \mathbf{j} \ \& \ \mathbf{j} \leq p \ \& \ \blacklozenge \mathbf{i} \leq q \ \& \ \mu^* Y.(\Diamond(p \wedge q) \wedge \Box Y) \leq \mathbf{m} \Rightarrow \mathbf{i} \leq \mathbf{m}.$$

This is now ready for the Ackermann rule (RA) to be applied twice to eliminate both p and q , producing

$$\mathbf{i} \leq \Diamond \mathbf{j} \ \& \ \mu^* Y.(\Diamond(\mathbf{j} \wedge \blacklozenge \mathbf{i}) \wedge \Box Y) \leq \mathbf{m} \Rightarrow \mathbf{i} \leq \mathbf{m}.$$

Appendix: Algebraic properties of additional operations on perfect distributive lattices

Let \mathbf{C} be a perfect distributive lattice. The map $\kappa: J^\infty(\mathbf{C}) \rightarrow M^\infty(\mathbf{C})$ is defined as by $\kappa(x) = \bigvee(C \setminus \uparrow x)$. This map is an order isomorphism between $J^\infty(\mathbf{C})$ and $M^\infty(\mathbf{C})$ with κ^{-1} defined for $m \in M^\infty(\mathbf{C})$ by $\kappa^{-1}(m) = \bigwedge(C \setminus \downarrow m)$. Note that if \mathbf{C} is an atomic Boolean algebra, then κ is simply the negation \neg .

The results at the end of this section are algebraic versions of results presented in Appendix A of [6]. We acknowledge that some of these results have been shown independently by Zhao [16].

Lemma .3. *Let \mathbf{C} be a completely distributive complete lattice. For any $x \in J^\infty(\mathbf{C})$, $m \in M^\infty(\mathbf{C})$ and any $c \in C$*

1. $x \not\leq \kappa(x)$;
2. $\kappa^{-1}(m) \not\leq m$;
3. $x \not\leq c$ if and only if $c \leq \kappa(x)$;
4. $c \not\leq m$ if and only if $\kappa^{-1}(m) \leq c$.

Proof. We prove (1) and (3). Suppose that $x \leq \kappa(x) = \bigvee(C \setminus \uparrow x)$. Since x is completely join-prime we get that there exists $p \in (C \setminus \uparrow x)$ such that $x \leq p$, a contradiction. For (3), if $x \not\leq c$, then $c \in (C \setminus \uparrow x)$ and by the definition of κ we see that $c \leq \bigvee(C \setminus \uparrow x) = \kappa(x)$. Now suppose that $c \leq \kappa(x)$. If $x \leq c$ then by transitivity we would get $x \leq \kappa(x)$, contradicting (1). \square

We will often make use of the following approach in our proofs.

Lemma .4. *Let $p, q \in A^\delta$.*

1. If $(x \not\leq q \Rightarrow x \not\leq p)$ for all $x \in J^\infty(\mathbf{A}^\delta)$, then $p \leq q$;
2. If $(p \not\leq m \Rightarrow q \not\leq m)$ for all $m \in M^\infty(\mathbf{A}^\delta)$, then $p \leq q$.

Proof. 1. We use the fact $J^\infty(\mathbf{A}^\delta)$ is join-dense in \mathbf{A}^δ . We have $p = \bigvee \{x \in J^\infty(\mathbf{A}^\delta) \mid x \leq p\}$ and $q = \bigvee \{x \in J^\infty(\mathbf{A}^\delta) \mid x \leq q\}$. If $x \not\leq q \Rightarrow x \not\leq p$ then $\{x \in J^\infty(\mathbf{A}^\delta) \mid x \leq p\} \subseteq \{x \in J^\infty(\mathbf{A}^\delta) \mid x \leq q\}$ and hence $p \leq q$. Part 2 follows using the meet-density of $M^\infty(\mathbf{A}^\delta)$ in \mathbf{A}^δ . \square

Let f be any map $f: \mathbf{A} \rightarrow \mathbf{B}$ and let $c \in A^\delta$. The extensions f^σ and f^π of f are defined by:

$$f^\sigma(c) = \bigvee \left\{ \bigwedge \{f(a) \mid k \leq a \leq u\} \mid k \leq c \leq u, k \in \mathbb{K}(\mathbf{A}^\delta), u \in \mathbb{O}(\mathbf{A}^\delta) \right\}$$

and

$$f^\pi(c) = \bigwedge \left\{ \bigvee \{f(a) \mid k \leq a \leq u\} \mid k \leq c \leq u, k \in \mathbb{K}(\mathbf{A}^\delta), u \in \mathbb{O}(\mathbf{A}^\delta) \right\}.$$

The following result holds for general lattices.

Lemma .5. [8, Lemma 4.3] *Let \mathbf{L} and \mathbf{M} be lattices, and let $f: \mathbf{L} \rightarrow \mathbf{M}$ be an order-preserving map.*

1. $f^\sigma(k) = \bigwedge \{f(a) \mid a \in L \text{ and } k \leq a\}$ for all $k \in \mathbb{K}(\mathbf{L}^\delta)$.
2. $f^\pi(u) = \bigvee \{f(a) \mid a \in L \text{ and } a \leq u\}$ for all $u \in \mathbb{O}(\mathbf{L}^\delta)$.

From this point onwards we will be working with a distributive lattice \mathbf{A} and its canonical extension \mathbf{A}^δ .

Lemma .6. *Let $f: \mathbf{A} \rightarrow \mathbf{A}$ be an order-preserving map and Let $k \in \mathbb{K}(\mathbf{A}^\delta)$ and $u \in \mathbb{O}(\mathbf{A}^\delta)$. Then for any $c \in A^\delta$,*

1. if $c \not\leq f^\sigma(k)$, there exists $a \in A$ such that $k \leq a$ and $c \not\leq f^\sigma(a) = f(a)$;
2. if $f^\pi(u) \not\leq c$, there exists $a \in A$ such that $a \leq u$ and $f(a) = f^\pi(a) \not\leq c$;

Proof. The proofs of (1) and (2) follow from parts (1) and (2) respectively of Lemma .5. \square

The Lemma below is an algebraic version of Corollary A.4 from [6].

Lemma .7. *Consider \mathbf{A}^δ and let $k \in \mathbb{K}(\mathbf{A}^\delta)$ and $u \in \mathbb{O}(\mathbf{A}^\delta)$. Then*

1. $\Box k \in \mathbb{K}(\mathbf{A}^\delta)$;
2. $\Diamond u \in \mathbb{O}(\mathbf{A}^\delta)$;

Proof. 1. Since k is closed we have $\Box k = \Box(\bigwedge \{a \in A \mid k \leq a\})$. As \Box is completely meet-preserving we see that $\Box k = \bigwedge \{\Box a \mid a \in A, k \leq a\}$ which is a closed element of \mathbf{A}^δ since A is closed under \Box . Item (2) follows from the fact that \Diamond is completely join-preserving. \square

The following lemma, a modified version of [6, Lemma A.5], is an algebraic version of the Esakia lemma.

Lemma .8. *For every up-directed set U of open elements, and every down-directed set D of closed elements:*

1. $\Box(\bigvee U) = \bigvee \{\Box u \mid u \in U\}$;
2. $\Diamond(\bigwedge D) = \bigwedge \{\Diamond d \mid d \in D\}$.

Proof. 1. The fact that $\Box(\bigvee U) \geq \bigvee \{\Box u \mid u \in U\}$ follows from the fact that \Box is order-preserving. Now let $m \in M^\infty(\mathbf{A}^\delta)$ such that $\Box(\bigvee U) \not\leq m$. By Lemma .6(2), there exists $a \in A$ such that $a \leq \bigvee U$ and $\Box a \not\leq m$. By compactness, there exists a finite set $F \subseteq U$ such that $a \leq \bigvee F$. Since U is up-directed, there exists $v \in U$ such that $\bigvee F \leq v$ and hence $a \leq v$. By the fact that \Box is order-preserving, we get $\Box a \leq \Box v$ and hence $\Box v \not\leq m$. Thus $\bigvee \{\Box u \mid u \in U\} \not\leq m$ and by Lemma .4(2) we have that $\Box(\bigvee U) \leq \bigvee \{\Box u \mid u \in U\}$.

2. The inequality $\Diamond(\bigwedge D) \leq \bigwedge \{\Diamond d \mid d \in D\}$ follows from the fact that \Diamond is order-preserving. Now suppose that $x \in J^\infty(\mathbf{A}^\delta)$ such that $x \not\leq \Diamond(\bigwedge D)$. Since $\bigwedge D$ is closed, we can use Lemma .6(1) to get $a \in A$ such that $\bigwedge D \leq a$ and $x \not\leq \Diamond(a)$. Again using the fact that $\bigwedge D$ is closed, we use the compactness of \mathbf{A}^δ to get a finite subset $F \subseteq D$ such that $\bigwedge F \leq a$. Since D is down-directed, the set F has a lower bound $e \in D$. Now $x \not\leq \Diamond(e)$ and hence $x \not\leq \bigwedge \{\Diamond d \mid d \in D\}$. Using Lemma .4(1) now gives us the required inequality. \square

Corollary .9. *Let $u \in \mathbb{O}(\mathbf{A}^\delta)$ and $k \in \mathbb{K}(\mathbf{A}^\delta)$. Then*

1. $\Box u \in \mathbb{O}(\mathbf{A}^\delta)$;
2. $\Diamond k \in \mathbb{K}(\mathbf{A}^\delta)$.

Proof. We show (1). The set $\{a \in A \mid a \leq u\}$ is up-directed and hence Lemma .8(1) gives us the equality $\Box u = \Box(\bigvee \{a \in A \mid a \leq u\}) = \bigvee \{\Box a \mid a \in A, a \leq u\}$. Thus $\Box u$ is open. Part (2) follows similarly using the corresponding statement from Lemma .8. \square

The result below is an Esakia-type lemma for the adjoint operations and the implications. Parts (1) and (2) are adaptations of Lemma A.7 from [6].

Lemma .10. *Let $U \subseteq \mathbb{O}(\mathbf{A}^\delta)$ be an up-directed set and let $D \subseteq \mathbb{K}(\mathbf{A}^\delta)$ be a down-directed set. Then*

1. $\blacksquare(\bigvee U) = \bigvee \{\blacksquare u \mid u \in U\}$;
2. $\blacklozenge(\bigwedge D) = \bigwedge \{\blacklozenge d \mid d \in D\}$;
3. $\bigwedge D - \bigvee U = \bigwedge \{d - u \mid d \in D, u \in U\}$;
4. $\bigwedge D \rightarrow \bigvee U = \bigvee \{d \rightarrow u \mid d \in D, u \in U\}$.

Proof. 1. The fact that $\bigvee \{\blacksquare u \mid u \in U\} \leq \blacksquare(\bigvee U)$ follows from \blacksquare being order-preserving. Now suppose that $\blacksquare(\bigvee U) \not\leq m$. Thus $\kappa^{-1}(m) \leq \blacksquare(\bigvee U)$ and by the adjunction, $\Diamond(\kappa^{-1}(m)) \leq \bigvee U$. By Corollary .9(2) and the fact that $\kappa^{-1}(m)$ is closed, we can apply compactness to get a finite subset $F \subseteq U$ such that $\Diamond(\kappa^{-1}(m)) \leq \bigvee F$. Now since U is up-directed, there exists $v \in U$ such that $\Diamond(\kappa^{-1}(m)) \leq v$ and hence $\kappa^{-1}(m) \leq \blacksquare v$ and, furthermore, $\blacksquare v \not\leq m$. Thus $\bigvee \{\blacksquare u \mid u \in U\} \not\leq m$ and by Lemma .4(2) we have that $\blacksquare(\bigvee U) \leq \bigvee \{\blacksquare u \mid u \in U\}$.

2. Since \blacklozenge is order-preserving, we have that $\blacklozenge(\bigwedge D) \leq \bigwedge \{\blacklozenge d \mid d \in D\}$. Now suppose that $x \not\leq \blacklozenge(\bigwedge D)$ for $x \in J^\infty(\mathbf{A}^\delta)$. By Lemma .3(3) we have that $\blacklozenge(\bigwedge D) \leq \kappa(x)$ and by the adjunction we have $\bigwedge D \leq \Box(\kappa(x))$. By Corollary .9(1) and the fact that $\kappa(x)$ is open, we can apply compactness to get a finite subset $F \subseteq D$ such that $\bigwedge F \leq \Box(\kappa(x))$. Furthermore, since D is down-directed, there exists $e \in D$ with $e \leq \bigwedge F$. Hence $e \leq \Box(\kappa(x))$ and $\blacklozenge e \leq \kappa(x)$ and so $x \not\leq \blacklozenge(e)$. Thus x is not a lower bound for $\{\blacklozenge d \mid d \in D\}$ and so by Lemma .4(1) we have the required inequality.

3. For any $d \in D$ we have $\bigwedge D \leq d$ and hence $\bigwedge D - \bigvee U \leq d - \bigvee U$. For any $u \in U$ we have $u \leq \bigvee U$ and hence for arbitrary d we have $d - \bigvee U \leq d - u$. Hence $\bigwedge D - \bigvee U \leq \bigwedge \{d - u \mid d \in D, u \in U\}$. For the opposite inequality, note that by density and the definition of $-$ we have

$$\bigwedge D - \bigvee U = \bigwedge \{w \in \mathbb{O}(\mathbf{A}^\delta) \mid \bigwedge D - \bigvee U \leq w\} = \bigwedge \{w \in \mathbb{O}(\mathbf{A}^\delta) \mid \bigwedge D \leq \bigvee U \vee w\}.$$

Now we see that $\bigvee U \vee w = \bigvee \{u \vee w \mid u \in U\}$. Since $w \in \mathbb{O}(\mathbf{A}^\delta)$, the set $\{u \vee w \mid u \in U\}$ is a set of open elements. Thus by compactness, for every $w \in \mathbb{O}(\mathbf{A}^\delta)$ such that $\bigwedge D \leq \bigvee U \vee w$, there exists some d_w, u_w such

that $d_w \leq u_w \vee w$, or, equivalently, $d_w - u_w \leq w$. This gives us $\bigwedge \{w \in \mathbb{O}(\mathbf{A}^\delta) \mid \exists d \in D, u \in U \text{ such that } d - u \leq w\} \geq \bigwedge \{d - u \mid d \in D, u \in U\}$, and hence $\bigwedge D - \bigvee U \geq \bigwedge \{d - u \mid d \in D, u \in U\}$.

4. For all $d \in D$ and $u \in U$ we have $\bigwedge D \leq d$ and $u \leq \bigvee U$ and hence $d \rightarrow u \leq d \rightarrow \bigvee U \leq \bigwedge D \rightarrow \bigvee U$. Thus $\bigvee \{d \rightarrow u \mid d \in D, u \in U\} \leq \bigwedge D \rightarrow \bigvee U$. For the reverse inequality, first observe that by density

$$\bigwedge D \rightarrow \bigvee U = \bigvee \{k \in \mathbb{K}(\mathbf{A}^\delta) \mid k \leq \bigwedge D \rightarrow \bigvee U\} = \bigvee \{k \in \mathbb{K}(\mathbf{A}^\delta) \mid k \wedge \bigwedge D \leq \bigvee U\}.$$

Using compactness we get $\{k \in \mathbb{K}(\mathbf{A}^\delta) \mid k \wedge \bigwedge D \leq \bigvee U\} = \{k \in \mathbb{K}(\mathbf{A}^\delta) \mid \exists d \in D, u \in U \text{ s.t. } k \wedge d \leq u\}$. Thus $\bigwedge D \rightarrow \bigvee U = \bigvee \{k \in \mathbb{K}(\mathbf{A}^\delta) \mid \exists d \in D, u \in U \text{ s.t. } k \leq d \rightarrow u\} \leq \bigvee \{d \rightarrow u \mid d \in D, u \in U\}$. \square

The next lemma is an adaptation of Proposition A.8 from [6].

Lemma .11. *Let $k \in \mathbb{K}(\mathbf{A}^\delta)$ and $u \in \mathbb{O}(\mathbf{A}^\delta)$. Then*

1. $\blacksquare u \in \mathbb{O}(\mathbf{A}^\delta)$;
2. $\blacklozenge k \in \mathbb{K}(\mathbf{A}^\delta)$;
3. $k \rightarrow u \in \mathbb{O}(\mathbf{A}^\delta)$;
4. $k - u \in \mathbb{K}(\mathbf{A}^\delta)$.

Proof. 1. It is clear that $\bigvee \{a \in A \mid a \leq \blacksquare u\} \leq \blacksquare u$. Suppose $x \in J^\infty(\mathbf{A}^\delta)$ and that $x \leq \blacksquare u$. By the adjunction this gives $\diamond x \leq u$. Since $x \in \mathbb{K}(\mathbf{A}^\delta)$ we have that $\diamond x = \diamond(\bigwedge \{a \in A \mid x \leq a\}) \leq u$. Since $\{a \in A \mid x \leq a\}$ is down-directed, by Lemma .8(2) we have that $\diamond x = \bigwedge \{\diamond a \mid x \leq a\} \leq u$. Since $\diamond x$ is closed, we can use compactness to get a finite set $\{a_i\}_{i=1}^n$ such that $\bigwedge \{\diamond a_i \mid 1 \leq i \leq n\} \leq u$. Since \diamond is order-preserving, we have $\diamond(\bigwedge \{a_i \mid 1 \leq i \leq n\}) \leq \bigwedge \{\diamond a_i \mid 1 \leq i \leq n\} \leq u$. Now $b = \bigwedge \{a_i \mid 1 \leq i \leq n\} \in A$ and $x \leq b$. Thus we have $b \in A$ such that $\diamond b \leq u$ (and hence $b \leq \blacksquare u$) with $x \leq b$. This gives us that $\blacksquare u \leq \bigvee \{a \in A \mid a \leq \blacksquare u\}$.

2. Clearly $\blacklozenge k \leq \bigwedge \{a \in A \mid \blacklozenge k \leq a\}$. Suppose $m \in M^\infty(\mathbf{A}^\delta)$ and that $\blacklozenge k \leq m$. By the adjunction, we have $k \leq \Box m$, which, by Lemma .8(1) gives us $k \leq \bigvee \{\Box a \mid a \in A \text{ and } a \leq m\} = \Box m$. By compactness we can get a finite subset $\{\Box a_i\}_{i=1}^n$ such that $k \leq \bigvee \{\Box a_i \mid 1 \leq i \leq n\}$ and since \Box is order-preserving, we have that $k \leq \Box \bigvee \{a_i \mid 1 \leq i \leq n\}$. Clearly $b = \bigvee a_i \in A$ and $b \leq m$. Thus there exists $b \in A$ such that $k \leq \Box b$, and hence $\blacklozenge k \leq b$ with $b \leq m$. This implies that $\bigwedge \{a \in A \mid \blacklozenge k \leq a\} \leq \blacklozenge k$.

3. By the join-density of $\mathbb{K}(\mathbf{A}^\delta)$ we get $k \rightarrow u = \bigvee \{c \in \mathbb{K}(\mathbf{A}^\delta) \mid c \leq k \rightarrow u\} = \bigvee \{c \in \mathbb{K}(\mathbf{A}^\delta) \mid c \wedge k \leq u\}$. Since every $a \in A$ is an element of $\mathbb{K}(\mathbf{A}^\delta)$ we have $\bigvee \{a \in A \mid a \leq k \rightarrow u\} = \bigvee \{a \in A \mid a \wedge k \leq u\} \leq k \rightarrow u$. Now suppose that $c \in \mathbb{K}(\mathbf{A}^\delta)$ and that $c \wedge k \leq u$. Now $c = \bigwedge \{b_i \mid b_i \in A, c \leq b_i\}$ and $c \wedge k = \bigwedge \{b_i \wedge k \mid c \leq b_i, b_i \in A\} \in \mathbb{K}(\mathbf{A}^\delta)$. By compactness there exists a finite subset of the b_i such that $\bigwedge_{j=1}^n (b_j \wedge k) = (\bigwedge_{j=1}^n b_j) \wedge k \leq u$. Now $\bigwedge_{j=1}^n b_j = a_c \in A$ and so for every $c \in \mathbb{K}(\mathbf{A}^\delta)$ such that $c \wedge k \leq u$, there exists $a_c \in A$ such that $c \leq a_c$ and $a_c \wedge k \leq u$. Thus $\bigvee \{a \in A \mid a \wedge k \leq u\} = \bigvee \{a \in A \mid a \leq k \rightarrow u\} = k \rightarrow u$ and so $k \rightarrow u \in \mathbb{O}(\mathbf{A}^\delta)$.

4. Observe that $k - u = \bigwedge \{w \in \mathbb{O}(\mathbf{A}^\delta) \mid k - u \leq w\} = \bigwedge \{w \in \mathbb{O}(\mathbf{A}^\delta) \mid k \leq w \vee u\}$. Since each $a \in A$ is an element of $\mathbb{O}(\mathbf{A}^\delta)$ we have that $\bigwedge \{a \in A \mid k - u \leq a\} = \bigwedge \{a \in A \mid k \leq u \vee a\} \geq k - u$. Let $w \in \mathbb{O}(\mathbf{A}^\delta)$ such that $k \leq u \vee w$. Now $w = \bigvee \{b_i \in A \mid b_i \leq w\}$ and $u \vee w$ is open. By compactness there exists a finite subset of the b_i such that $k \leq \bigvee_{j=1}^m (u \vee b_j) = (\bigvee_{j=1}^m b_j) \vee u$. Now $\bigvee_{j=1}^m b_j = a_w \in A$ and $k \leq a_w \vee u$ and $a_w \leq w$. Therefore $\bigwedge \{a \in A \mid k \leq u \vee a\} \leq \bigwedge \{w \in \mathbb{O}(\mathbf{A}^\delta) \mid k \leq u \vee w\} = k - u$ and so $k - u \in \mathbb{K}(\mathbf{A}^\delta)$. \square

References

- [1] S. AMBLER, M. Z. KWIATKOWSKA, AND N. MEASOR, *Duality and the completeness of the modal mu-calculus*, Theor. Comput. Sci., 151 (1995), pp. 3–27.
- [2] N. BEZHANISHVILI AND I. HODKINSON, *Preservation of Sahlqvist fixed point equations in completions of relativized fixed point Boolean algebras with operators*, Algebra Universalis, 68 (2012), pp. 43–56.

- [3] ———, *Sahlqvist theorem for modal fixed point logic*, Theor. Comput. Sci., 424 (2012), pp. 1–19.
- [4] J. BRADFIELD AND C. STIRLING, *Modal μ -calculi*, in Handbook of Modal Logic, P. Blackburn, J. van Benthem, and F. Wolter, eds., Elsevier, 2006, pp. 721–756.
- [5] W. CONRADIE, Y. FOMATATI, A. PALMIGIANO, AND S. SOURABH, *Algorithmic correspondence for intuitionistic μ -calculus*, Theor. Comput. Sci. to appear.
- [6] W. CONRADIE AND A. PALMIGIANO, *Algorithmic correspondence and canonicity for distributive modal logic*, Ann. Pure Appl. Logic, 163 (2012), pp. 338–376.
- [7] G. FISCHER-SERVI, *Semantics for a class of intuitionistic modal calculi*, in Italian Studies in the Philosophy of Science, M. L. Dalla Chiara, ed., vol. 47 of Boston Studies in the Philosophy of Science, Springer Netherlands, 1981, pp. 59–72.
- [8] M. GEHRKE AND J. HARDING, *Bounded lattice expansions*, J. Algebra, 238 (2001), pp. 345–371.
- [9] M. GEHRKE, Y. NAGAHASHI, AND Y. VENEMA, *A Sahlqvist theorem for distributive modal logic*, Ann. Pure Appl. Logic, 131 (2005), pp. 65–102.
- [10] V. GORANKO AND D. VAKARELOV, *Elementary canonical formulae: extending Sahlqvist’s theorem*, Ann. Pure Appl. Logic, 141(1–2) (2006), pp. 180–217.
- [11] D. KOZEN, *Results on the propositional μ -calculus*, Theor. Comput. Sci., 27 (1983), pp. 333–353.
- [12] H. SAHLQVIST, *Correspondence and completeness in the first and second-order semantics for modal logic*, in Proceedings of the 3rd Scandinavian Logic Symposium, Uppsala 1973, S. Kanger, ed., Springer-Verlag, Amsterdam, 1975, pp. 110–143.
- [13] J. VAN BENTHEM, *Modal Logic and Classical Logic*, Bibliopolis, 1983.
- [14] J. VAN BENTHEM, *Correspondence theory*, in Handbook of Philosophical Logic, Vol. 3, D. Gabbay and F. Guenther, eds., Kluwer Academic Publishers, 2001, pp. 325–408.
- [15] F. WOLTER AND M. ZAKHARYASCHEV, *Intuitionistic modal logic*, in Logic and Foundations of Mathematics, A. Cantini, E. Casari, and P. Minari, eds., vol. 280 of Synthese Library, Springer Netherlands, 1999, pp. 227–238.
- [16] Z. ZHAO, *Algebraic Canonicity in Non-Classical Logics*, 2013. MSc Thesis, ILLC, University of Amsterdam, The Netherlands. Available at www.illc.uva.nl/Research/Reports/MoL-2013-10.text.pdf.